**Fuzzy Management Methods**
*Series Editors:* Andreas Meier · Witold Pedrycz · Edy Portmann

Miloš Švaňa
František Zapletal
Miroslav Hudec

# From Text to Understanding

## Using Fuzzy Sets to Analyse Free-Form Text Data

Springer

# Fuzzy Management Methods

**Series Editors**
Andreas Meier
*Fribourg, Switzerland*

Witold Pedrycz
*Edmonton, Canada*

Edy Portmann
*Bern, Switzerland*

With today's information overload, it has become increasingly difficult to analyze the huge amounts of data and to generate appropriate management decisions. Furthermore, the data are often imprecise and will include both quantitative and qualitative elements. For these reasons it is important to extend traditional decision making processes by adding intuitive reasoning, human subjectivity and imprecision. To deal with uncertainty, vagueness, and imprecision, Lotfi A. Zadeh introduced fuzzy sets and fuzzy logic. In this book series "Fuzzy Management Methods" fuzzy logic is applied to extend portfolio analysis, scoring methods, customer relationship management, performance measurement, web reputation, web analytics and controlling, community marketing and other business domains to improve managerial decisions. Thus, fuzzy logic can be seen as a management method where appropriate concepts, software tools and languages

build a powerful instrument for analyzing and controlling the business.

Miloš Švaňa, František Zapletal and Miroslav Hudec

# From Text to Understanding
## Using Fuzzy Sets to Analyse Free-Form Text Data

Miloš Švaňa
VSB – Technical University of Ostrava, Ostrava-Poruba, Czech Republic

František Zapletal
VSB – Technical University of Ostrava, Ostrava-Poruba, Czech Republic

Miroslav Hudec
VSB – Technical University of Ostrava, Ostrava-Poruba, Czech Republic

ISSN 2196-4130 e-ISSN 2196-4149
Fuzzy Management Methods
ISBN 978-3-032-00128-3 e-ISBN 978-3-032-00129-0
https://doi.org/10.1007/978-3-032-00129-0

© The Editor(s) (if applicable) and The Author(s), under exclusive license to Springer Nature Switzerland AG 2025

This work is subject to copyright. All rights are solely and exclusively licensed by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not

imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

*To my girlfriend Lenka, who supported and motivated me during hard times, and to all my teachers who helped turn a curious child into a doctor.*
*Miloš Švaňa*
*To my wife Katarína for her love and kind support, and kids František and Kristýna for teach me how to strengthen my patience and reshuffling my life preferences in a definitely better manner.*
*František Zapletal*
*To my parents, whose words of encouragement supported my journey to knowledge. To my wife for her incredible level of love and patience. To my children, who fill my heart with joy and mind with research questions.*
*Miroslav Hudec*

# Foreword

The quest for *smarter*, more responsive urban environments stands as one of the defining challenges of our contemporary era. As cities face an array of mounting complexities-including issues surrounding mobility, sustainability, and social cohesion—the vision of the smart city has emerged as an enticing beacon for future progress. However, this monograph thoughtfully exposes a deeper and critical democratic question that lurks beneath the polished surface of algorithmic governance: whose voices are shaping these urban futures?

This monograph represents a timely and methodical contribution to that ongoing discourse. It rigorously examines the prevailing imbalance within municipal decision-making processes, where the power and influence of technology vendors often eclipse the voices and needs of everyday citizens, in an environment where traditional survey methods and participatory forums frequently fall short-whether due to high costs, exclusionary practices, or ineffectiveness—the thesis boldly proposes a transformative alternative: tapping into the rich civic conversations that unfold within social media platforms. The author goes beyond merely advocating for the collection of more data; instead, the focus is on garnering better insights-insights that honor the diversity of opinion, the nuance of sentiment, and the localized contexts that shape civic discussions. At the heart of this work lies an innovative analytical framework that synthesizes natural language processing, fuzzy set theory, and design science methodology.

Through this interdisciplinary approach, the thesis seeks to address two critical gaps: the chronic underrepresentation of citizen perspectives in urban

governance and the tendency to oversimplify sentiments in existing analytical methodologies. More than a mere technical contribution, this research paves the way for ethical, inclusive, and explainable civic intelligence. It illustrates how digital artifacts, when thoughtfully designed, can cultivate not only smart cities but also wise and equitable urban environments.

Grounded in practical relevance, the case studies focused on Czech municipalities lend credence to the framework. At the same time, the robust theoretical scaffolding establishes a solid foundation for future inquiries across various fields, including urban studies, data science, and public administration. For scholars, policymakers, and technologists alike, this thesis serves as an urgent call to reimagine digital participation, not as an ancillary convenience added to governance but as an essential principle guiding the design of urban futures. By providing structure to complexity and amplifying the voices of the unheard, it illuminates a crucial pathway forward for cities that aspire not just to be efficient but also truly equitable.

**Luis Terán**
**July 2025**

# Preface

Many cities around the world are aspiring to become *smart*. A *smart city* deploys information technologies to make better decisions in areas like traffic and transportation, environmental protection, waste management, housing, or culture.

Municipal planning involves multiple stakeholders. Local government, companies, citizens, commuters, and environmental activists are among the most important. Some of these groups are more influential than others. Many smart city and urban planning experts criticize the disproportionately large involvement of companies whose business revolves around smart city solutions. These experts claim that decision-making at the municipal level is often predominantly based on the interplay between such companies and local governments.

In contrast, the opinions of average citizens often have little weight. Efforts to gather these opinions are often half-hearted; smart city events are often inaccessible to a layperson, and discussion forums are often held during working hours, which significantly limits attendance. Ignoring citizen opinions can lead to a mismatch of priorities between local governments and citizens.

Cities can have valid reasons for this perceived lack of interest. Traditional surveys can be costly and labor-intensive in complex citizens opinions evaluations. Other disadvantages include the inability to reliably generalize survey results, low response rates, no opportunity to ask follow-up questions, lack of explanation of the observed patterns and trends, question misinterpretation, lack of control over timeliness, truthfulness of the answers, or lack of knowledge on the part of the respondent.

Social networks are a popular tool for sharing opinions that can serve as an alternative or a complement to tools such as surveys or interviews. Because of their large user bases, they might contain more useful information than can be extracted by other means. Social network users are not limited to expressing their opinions through a predetermined set of questions and answers. They can post various types of content, including free-form text, images, videos, links, or even simple polls. They can also discuss unanticipated novel topics. But the lack of formal structure makes processing and interpreting this data difficult.

Social network data are commonly analyzed with natural language processing methods, such as sentiment analysis and topic modeling. Several commercial tools offer this type of analysis. However, most of these tools provide only surface-level information. For example, they often simplify sentiment evaluation to labeling documents as *positive*, *negative*, or *neutral*. Current solutions also often ignore opinion diversity and the aggregation of sentiment across documents. Commercial tools reduce this problem to counting positive, negative, and neutral documents. However, this information is difficult to interpret. The ratio of positive and negative documents can be understood as either average sentiment or a degree of diversity.

In short, this book addresses two issues:

- Citizen opinions are underrepresented in municipal decision-making.
- Methods for analyzing social media data often ignore opinion diversity.

We propose a novel framework to support municipal decision-making by analyzing social media data while capturing opinion diversity. The framework should help municipalities make better decisions by identifying topics discussed by social media users located in a given municipality, analyzing sentiment toward these topics

together with their diversity, and identifying controversial or negatively perceived topics.

Social media data and the proposed framework could represent a cheaper and more scalable alternative to other opinion collection tools, such as surveys or interviews. Information about diversity could significantly improve municipal decision-making. High diversity in opinions toward a topic could prompt a deeper investigation, which in turn would lead to bringing the real issues to the surface. This advantage would be lost if sentiment were aggregated by a simple method such as the arithmetic mean. For example, it would be impossible to distinguish between a slightly positive average sentiment resulting from all people having slightly positive opinions and a slightly positive average sentiment resulting from some people having very positive opinions and some people having slightly negative opinions.

We approach the problem of incorporating citizen opinion in municipal decision-making and the problem of dealing with uncertainty associated with these opinions by applying a methodology called Design Science Research (DSR). The following description is taken from [1].

DSR is focused on the *science of the artificial*—people-made artifacts such as algorithms, physical or abstract processes, or various physical devices. DSR uses artifact design as a tool to create new knowledge about how an artifact solving a specific problem can be built.

DSR differs from routine industrial design in its goal. The goal of DSR is to create new and true knowledge. In contrast, the goal of routine design could be summarized as finding a practical solution to a real-world problem. Routine design often reuses existing components or processes to reduce risk, while DSR reduces uncertainty by creating knowledge.

The DSR process consists of five main stages:

1. **Awareness of Problem** can come from many different sources, including new developments in industry, scientific fields of interest, or by studying an alien discipline and realizing that knowledge from this discipline can be applied in one's main field of research.

2. **Suggestion**. The output of this phase is a high-level tentative design that could solve a given problem (or the realization that no solution presents itself to the researcher).

3. **Development** involves the actual creation of the artifact. Many standard procedures built on top of existing knowledge can be used in this phase, since the artifact itself is only a research tool.

4. **Evaluation**. After the artifact is created, it is necessary to evaluate the hypotheses about its usefulness in solving a given problem, as well as the hypotheses about its behavior. If the artifact significantly deviates from expectations, its design must be modified.

5. **Conclusion**. A DSR project usually ends when an artifact is *good-enough*. The artifact does not have to represent a perfect solution to the problem at hand, as the goal of DSR is to create knowledge about how artifacts solving a given problem can be designed.

The book is organized as follows:

Chapter 1 introduces the context for the two problems stated above. It defines what a smart city is and lists the most important issues related to the *smartness* of cities. Then, the focus is narrowed down to municipal decision-making. We describe both the current state and the new emerging trends. Finally, we introduce the concept of

*cognitive cities* as a new approach to addressing some of the issues we identified.

Chapter 2 introduces fuzzy sets and fuzzy algebra as key methodological tools used to develop the aforementioned framework. We define several general concepts from fuzzy set theory and then focus on concepts relevant to the framework design, such as fuzzy set similarity or approaches to estimating the shapes of fuzzy sets.

The framework also utilizes natural language processing (NLP) techniques such as topic modeling and sentiment analysis. These techniques are themselves based on machine learning (ML) methods. Therefore, Chap. 3 provides a short theoretical overview of both fields.

We introduce the design of the framework in Chap. 4. We describe both the overall structure of the framework and justify our design decisions, including the choice of particular sentiment analysis and topic modeling methods. We also explain how we combine topic modeling, sentiment analysis, and fuzzy set theory to capture and express the information about average sentiment and sentiment diversity.

Chapter 5 summarizes real-world case studies and experiments performed to test the framework. This includes experiments on a synthetic dataset verifying that the framework behaves as intended, case studies on social media posts from two Czech cities, and a case study conducted on a well-known Sentiment140 dataset. We demonstrate how a decision-maker can derive practical insights from the results provided by the framework, we compare the framework with a naive approach to sentiment analysis, and empirically evaluate its computational complexity.

Finally, Chaps. 6 and 7 discuss limitations of the framework with a focus on possible biases, other potential

applications, sustainability implications, and future research opportunities.

---

## Reference

1. Vaishnavi V, Kuechler B, Petter S (2004) Design science research in information systems. http://www.desrist.org/design-research-in-information-systems/

**Miloš Švana**
**František Zapletal**
**Miroslav Hudec**
**Ostrava, Czech Republic**
**May 2025**

# Acknowledgments

This work is supported by the European Union under the project REFRESH —Research Excellence For Region Sustainability and High-tech Industries project number CZ.10.03.01/00/22_003/0000048 via the Operational Programme Just Transition as well as the SGS project no. SP2025/081 of the Ministry of Education, Youth and Sports of the Czech Republic.

Advice and suggestions on theoretical and practical implications related to our work were a valuable source for this research. In this direction, we express our deep appreciation to Edy Portmann, José María Alonso Moral, Daša Mokošová, Galena Pisoni, Jörg Osterrieder, Miljan Vučetić, Henri Pirkkalainen, and Hongxiu Li for their insightful ideas and challenging critical comments.

We are also grateful to the management of the Faculty of Economics, VŠB - Technical University of Ostrava, for providing us with space for research which enabled us to spend long hours in researching and writing the book, as well as to our colleagues and students in the Department of System Engineering and Informatics for tedious work in labeling sentiment data.

Our final thanks go to Springer for offering the opportunity to publish our research in their publishing house, and especially to Barbara Bethke and Parthiban Gujilan Kannan for cooperating during the entire publishing process.

# Acronyms

*API* Application Programming Interface

*BERT* Bidirectional Encoder Representations from Transformers

*CBOW* Continuous Bag of Words

*DBSCAN* Density-Based Space Clustering of Applications with Noise

*DMM* Dirichlet Multinomial Mixture

*DSR* Design Science Research

*GDP* Gross Domestic Product

*GPT* Generative Pre-trained Transformer

*ICT* Information and Communications Technology

*IoT* Internet of Things

*HDBSCAN* Hierarchical Density-Based Spatial Clustering of Applications with Noise

*LDA* Latent Dirichlet Allocation

*LSA* Latent Semantic Analysis

*ML* Machine Learning

*NFM* Non-negative Matrix Factorization

*NLP* Natural Language Processing

*TF-IDF* Term Frequency-Inverse Document Frequency

*TFN* Triangular Fuzzy Number

*UMAP* Uniform Manifold Approximation and Projection

*VADER* Valence Aware Dictionary and sEntiment Reasoner

# Contents


1 Citizen Engagement in Smart Cities Municipal Decision-Making

- 1.1 What Is a Smart City?
- 1.2 Components of a Smart City
- 1.3 Models of Smart City Development
- 1.4 Current Issues in Smartness of Cities
- 1.5 Models of Decision-Making and Citizen Engagement
- 1.6 New Trends in Citizen Engagement
- 1.7 From Smart to Cognitive Cities
- References

**Part I Methodology**

2 Fuzzy Sets and Fuzzy Algebra

- 2.1 Introduction to Fuzzy Mathematics
  - 2.1.1 Fuzzy Sets, Their Description, and Properties
- 2.2 Fundamentals of Fuzzy Algebra
  - 2.2.1 Similarity of Fuzzy Sets
- 2.3 Linguistic Variables
- 2.4 Estimating Membership Functions
  - 2.4.1 Data-Driven Approaches
  - 2.4.2 User-Driven Approaches
- References

3 Machine Learning and Natural Language Processing
3.1 Classification of Machine Learning Approaches
3.1.1 Supervised Learning
3.1.2 Unsupervised Learning
3.1.3 Reinforcement Learning
3.1.4 Hybrid Approaches
3.2 Machine Learning Methods
3.2.1 Naive Bayes
3.2.2 HDBSCAN
3.2.3 Neural Networks and Transformers
3.3 Natural Language Processing
3.4 Topic Modeling
3.4.1 Topic Model Evaluation
3.4.2 Choosing an Appropriate Model
3.4.3 Latent Dirichlet Allocation
3.4.4 BERTopic
3.5 Sentiment Analysis
References

**Part II Framework Design and Numerical Experiments**

4 Framework Design
4.1 Introduction to Framework Design
4.2 Evaluation of Sentiment Analysis Methods
4.3 Evaluation of Topic Modeling Methods

4.4 Fuzzy Sentiment Aggregation
4.5 Levels of Positive and Negative Opinions
References
5 Numerical Experiments
5.1 Overview
5.1.1 Synthetic Dataset
5.2 Case Study of Ostrava and Brno
5.2.1 Preprocessing
5.2.2 Basic Data Analysis
5.2.3 Tweets from Ostrava
5.2.4 Tweets from Brno
5.2.5 Comparison with a Naive Method
5.2.6 Likes and Shares
5.2.7 Web Application
5.3 Experiments with the Sentiment140 Dataset
5.3.1 Comparison with a Naive Approach
5.3.2 Processing Time
References
Part III Conclusions and Recommendations
6 Discussion and Recommendations
6.1 Biases
6.2 Determining the Shapes of Positive and Negative Opinions
6.3 Other Methodological Issues

6.4 Applications
References
7 Conclusions
7.1 Ethics and Sustainability
7.2 Future Research Opportunities
7.3 Concluding Remarks
References
Index

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2025
M. Švaňa et al., *From Text to Understanding*, Fuzzy Management Methods
https://doi.org/10.1007/978-3-032-00129-0_1

# 1. Citizen Engagement in Smart Cities Municipal Decision-Making

Miloš Švaňa[1] ✉, František Zapletal[2] ✉ and Miroslav Hudec[3] ✉

(1) VSB – Technical University of Ostrava, Ostrava-Poruba, Czech Republic
(2) VSB – Technical University of Ostrava, Ostrava-Poruba, Czech Republic
(3) VSB – Technical University of Ostrava, Ostrava-Poruba, Czech Republic

✉ **Miloš Švaňa (Corresponding author)**
**Email:** milos.svana@gmail.com

✉ **František Zapletal (Corresponding author)**
**Email:** frantisek.zapletal@vsb.cz

✉ **Miroslav Hudec (Corresponding author)**
**Email:** miroslav.hudec@vsb.cz

## 1.1 What Is a Smart City?

The urban system has become the predominant habitat of the global population. Cities are described as “the highest forms of economic and sociocultural achievements of

human civilization” and as “engines of economic growth” [29], where it was claimed that at the time of writing approximately 73% of Europeans lived in cities and that in developing parts of the world, such as Asia and Africa, this ratio grows rapidly. Therefore, it is not surprising that about 80% of the world’s GDP is generated by urban areas [4].

Smart cities are characterized by an extensive deployment of information and communication technologies in various areas of urban life, including optimizing transportation, accommodation, environment, or waste management. Internet of Things (IoT) [20] plays an important role in this direction. The next observation is that urban systems are not merely technical artifacts [14]; rather, they are complex sociotechnical systems in which citizens should actively participate in ways suitable for all categories of citizens.

The term *smart city* was first introduced by Gibson, Kozmetsky, and Smilor [18, 30]. There is no agreed definition of a smart city. Kumar [29] and Koutra [30] summarize several definitions proposed by different authors, for example:

**Definition 1.1** A smart city is characterized by the use of smart computing technologies to make the critical infrastructure components and service of a city including administration, education, health care, public safety, real estate, transportation, and utilities more intelligent, interconnected, and efficient [48].

**Definition 1.2** A smart city is a city that performs well in a forward-looking way in economy, people, governance, mobility, environment, and living built on the smart combination of endowments and activities of self-decisive, independent, and aware citizens [19].

**Definition 1.3** Smart cities are cities that use information and communication technologies in order to increase the quality of life of their inhabitants while contributing to sustainable development [9].

**Definition 1.4** A smart city is a city that monitors and integrates the conditions of all its critical infrastructure, including roads, bridges, tunnels, rails, subways, airports, seaports, communications, water, power, and even major buildings, can better optimize its resources, plan its preventive maintenance activities, and monitor security aspects while maximizing services to its citizens [37].

The concept of smart cities is linked to Goal 11 of the United Nations' Sustainable Development Goals: "Make cities and human settlements inclusive, safe, resilient, and sustainable." Three elements play a role in achieving sustainable urban growth: (1) national urban policies, (2) laws, institutions, and systems of governance, and (3) urban economy [28].

Smart city research is very active and is cross-disciplinary and cross-sectional [46]. It involves participants from government, educational, and research institutions, as well as industry. Moreover, it is a worldwide phenomenon. The interest in smart cities can be seen both in the developed world, including the United Kingdom [46] or South Korea [50], and in developing countries such as India [21] or Indonesia [4].

---

## 1.2 Components of a Smart City

The idea of a smart city can be further divided into six components: (1) Smart People, (2) Smart Economy, (3) Smart Mobility, (4) Smart Environment, (5) Smart Living, and (6) Smart Governance [10, 28]. They are depicted in Fig. 1.1

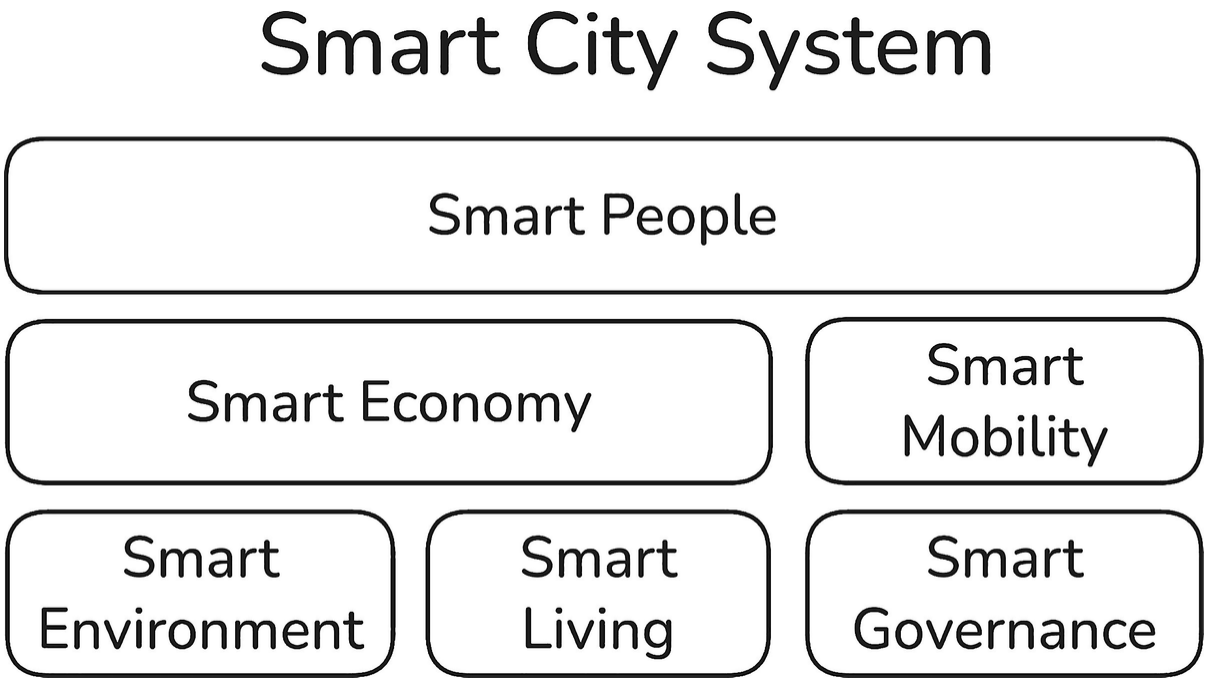


***Fig. 1.1*** Building blocks of a smart city [28]

**Smart People** should "excel in what they do professionally" and "have a high level of Human Development Index." They are lifelong learners, flexible, resilient to changing circumstances, creative, cosmopolitan, open-minded, and actively involved in municipal decision-making. Smart cities should support the "smartness" of people by "integrating its universities and colleges into all aspects of city life," by attracting knowledge workers and other human capital, and by maintaining a high graduate enrollment ratio [29].

A city with a **Smart Economy** understands its "economic DNA," is driven by innovation, values creativity, has entrepreneurial leadership, and offers diverse economic opportunities to its citizens. It is prepared for the challenges and opportunities. Sustainable growth and the management of natural resources (unfortunately, often just empty words) are crucial to their strategy.

**Smart Mobility** means focusing on walking, cycling, and balanced transportation options. Solutions supporting this goal include vibrant streets, high-density living, high-

quality and fast mass transportation, or efficient traffic management.

**Smart Environment** covers both natural protection and efficient management of natural resources, such as water. Smart cities should offer public green spaces for various outdoor activities and survival for wild animals such as birds. Cities must efficiently manage waste and energy, focus on low carbon and renewable energy solutions, and prepare themselves for the negative impacts of the serious effects of climate change (people do not want to sacrifice their comfort) in the future.

In the area of **Smart Living**, the city should focus not only on its culture but also on safety and security. It should offer a vibrant downtown, high-quality public services and amenities, support art and various cultural events, and enrich the aesthetics of daily life.

Finally, **Smart Government** should be accountable, responsive, and transparent. It should use primarily input from citizens and also big data and decision support systems to make decisions that lead to efficient resource management. It should implement solutions that support citizen participation in decision-making.

This book focuses on the last component of the smart city. In this component, social media sentiment analysis might, among others, help local governments to be more responsive to the needs of their citizens. It can also be seen as a form of indirect citizen participation in municipal decision-making.

---

## 1.3 Models of Smart City Development

Three basic models of smart city development include [22]:

1. Smart-from-the-start-cities

2. Retrofitted smart cities

3. Social cities

**Smart-from-the-start** cities are built from scratch with the goal of creating a perfect smart city. Such initiatives can be found all over the world. They are often located in free economic zones, where many laws of the host country do not apply. The problem with this model is that smart-from-the-start cities are not very successful. People do not want to live in them, and many describe them as “cold” or “soulless.” Moreover, the construction of smart-from-the-start cities often leads to forced relocation and land buyouts at unfair prices. An example is the experiment [22], which ended as a weakly occupied city, because the inhabitants felt that a lot of sensitive data were collected for optimization purposes and therefore citizens felt like the walking IoT.

**Retrofitted smart cities** are existing cities that gradually implement new smart city technologies. Halegoua [22] sees practical realizations of this model in a rather negative light: Related projects are often dominated by a few big corporations that often create their own demand by providing consulting services or even initial financing. These companies sell all-in-one smart city packages. They make smart cities a commodity while ignoring social aspects [14]. Retrofitted smart cities often lack a coherent vision, while focus only on one or two selected areas, such as water management or transportation.

Both smart-from-the-start and retrofitted smart cities are centered on the idea of economic growth and ensuring efficiency. Optimization methods have been applied to find the best solution for a variety of problems. Clearly, this way is significant for improving living conditions in urban

systems, but it is simply not sufficient. One should go beyond pure technological concepts and include social aspects such as citizen participation and informing about development in an understandable way, because urban systems should not be merely technical artifacts [25]. For example, work [34] has merged these two aspects in a technological solution based on fuzzy logic to provide efficient last-mile delivery respecting the privacy of citizens.

In contrast, the main goal of **social cities** is to create places where people actually want to live. To make a city liveable, factors such as social life, culture, and communities should be a priority. Social cities focus on citizen participation. Instead of relying on big companies, they let people define and solve their own problems, for example, through various initiatives.

Some authors take a more normative approach to smart city development. For example, [46] describes a smart city development framework consisting of five elements depicted in Fig. 1.2:

1. Identification of real needs

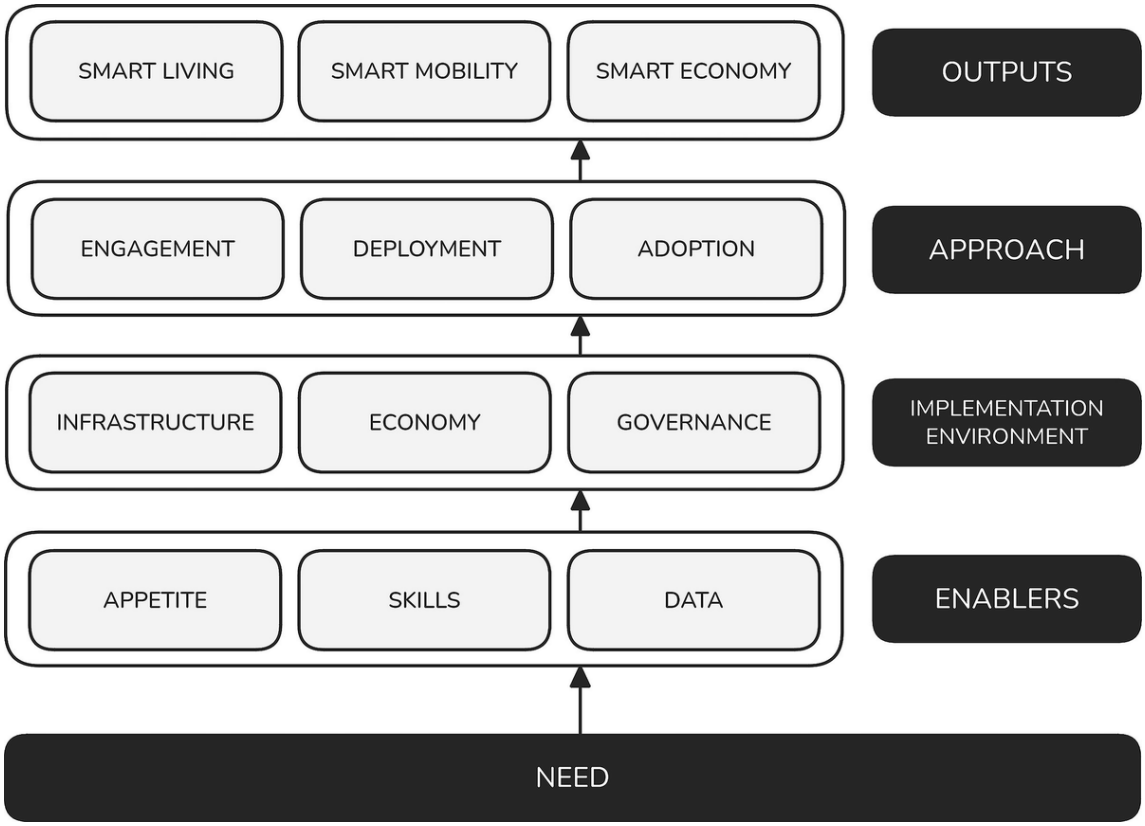


***Fig. 1.2*** Smart city development framework proposed by [46]

2. Enablers such as skilled workforce or citizen appetite

3. Implementation environment, such as current infrastructure and institutions

4. Approach to engagement and implementation

5. Outputs

The authors emphasize that all components are mandatory. They must be approached as an integrated whole. However, when taken as a research objective, it might still make sense to analyze each component individually. This work is tied to component no. 1, the key idea being the utilization of social media data to provide a more faithful representation of citizen needs.

Around the world, researchers and practitioners focus their work on improving living conditions in cities, i.e., making cities smarter. It is a difficult task to encompass optimization and citizen needs in an uncertain and dynamic environment. This book focuses on issues at the intersection of retrofitted and social city models. Its goal is to support the inclusion of the opinions and priorities of citizens in municipal decision-making. At the same time, the proposed framework is intended to be deployed in existing cities.

---

## 1.4 Current Issues in Smartness of Cities

Although the concept of smart cities has existed for some time, many issues related to “smartness” still exist. Many smart cities are not really smart, even if they claim so. Real smartness means addressing real problems [46]. It also means a high level of interconnectedness and interactivity between different subsystems and the participation of citizens. However, this aspect is often missing [50].

Cities also lack a good long-term vision of what they want to achieve [46]. Moreover, urban development is often slowed by a lack of skilled workers [29]. Cities in some regions, such as South Korea, underestimate the protection of nature [50]. The results of the survey conducted in Czechia [16] reveal many issues related to Prague’s urban development, including a lack of long-term planning, a gap between the policies on paper and their implementation, or an emphasis on cars at the expense of other forms of transportation. In a more general sense, the mismanagement of resources, including data and local communities, can result in a failed implementation of a smart city [4].

Shishanina and Sidorov [40] explore issues in urban development related to bureaucracy and management in Russia. They were able to identify several important problems:

- Nonsystemic structure of the strategic development process
- Predominant reliance on administrative methods of management
- Insufficient methodological support of local initiatives from the regional-level government
- Skill gap in strategic and project management in public officials at various levels
- Heavy financial dependence, since municipalities usually have no additional sources of income
- A lack of an effective mechanism for collaboration of local authorities with the local community, business, and other stakeholders in the development

A comprehensive review of the challenges faced by smart city initiatives is provided in [36]. That work starts by highlighting that all cities are limited by available financial, material, and human resources. Next, they describe several specific issues and categorize them into technological and nontechnological. The most important technological issues include privacy, security, and interoperability between different systems. By far, the most pressing nontechnological challenge is collaboration—both internal (between municipal departments) and external (municipal authorities and other stakeholders). Another important nontechnological challenge is financing, as smart city projects are risky, costly, and profitable only in the long term, but the question is whether profitability is more relevant than citizen comfort. A complete list of the issues identified by the authors can be found in Table 1.1.

***Table 1.1*** Smart city issues identified by Pierce and Andersson [36]

| Nontechnical challenges | Technical challenges |
|---|---|
| **Collaboration** | **Security** |
| **Weak collaborative engagement with external stakeholders** | Big data in public clouds brings with it issues relating to security |
| **Weak collaborative engagement with internal stakeholders** | Challenges of smart city infrastructure integration |
| **The lack of an aligned vision of the city development with external and internal stakeholders** | Outsourcing of power and control to private sector providers |
| **Issues concerning institutional resistance** | |
| **Financial** | **Interoperability** |
| **Limited funds for smart city initiatives** | Issues of service interoperability |
| **Long-term delay before reaching maturity and profitability** | Smart city infrastructure integration: system interoperability |
| **Large up-front investment** | |
| **Difficult to monetize on smart city investments** | |
| **Governance** | **Privacy** |
| **Control creep, the hollowing out of state provided services, widening inequalities and dispossession of land and livelihoods** | How to make sense and best use of such “big data,” while preserving citizens’ privacy and data security |
| **Formulating the right governance structure given organizational barriers to supporting smart city initiatives** | Key issue relates to the privacy of the information monitored by sensors and to the implications that the violation of this principle can have on citizens’ routines and habits in case of malicious or unintentional data exposure |
| | Smart city technologies may encourage increased surveillance |
| **Contextual** | |

| Nontechnical challenges | Technical challenges |
|---|---|
| **Influence of geographical variables and exportation of best practices may not occur easily** | |
| **Difficult to extend local imitations to multicity projects due to the localized character of initiatives** | |
| **The use of canonical examples and one-size-fits-all narratives** | |
| **Political** | |
| **Political uncertainties hampering public and private investments** | |
| **Political hyper-activism** | |

Several authors discuss potential solutions to these issues. For example, Koutra et al. [30] suggest that smart city planning should be based on a bottom-up approach, starting at the level of individual districts. There are also many opportunities to integrate smart city development into higher education [8].

The issue this book is addressing relates to the last two of the six components constituting a smart city described above: smart people and smart governance. Smart city development is focused on the technological side and social aspects, while citizen opinions are often ignored [22]. This is consistent with some authors who describe smart people as the main component of smart city development, while in reality too much emphasis is placed on financing various technological projects [28]. Truly smart cities cannot be created by “pouring” public or private money into city development, but rather by a deliberate decision of people living in a given city to implement some aspects of a smart city.

In the work [49], the author asks the question *whose smartness and whose cities*? Referring to many other authors, the work argues that citizen participation is often mentioned in smart city policy documentation but neglected in practice. One of the reasons is the lack of tools that would facilitate citizen participation. Moreover, the current conception of smart cities requires *smart citizens* who are willing to adapt and incorporate various technologies into their lives. This can lead to discrimination of certain groups, for example, those who are technologically illiterate and reluctance of citizens that are not keen to be considered as a walking IoT (when the solution expects heavy use of new technologies). This view is consistent with [46] claiming that citizen engagement is a key factor in determining what issues should be addressed by smart city projects.

Moreover, multiple studies criticize the top-down approach to urban planning and state that the focus on efficiency and sustainability may sometimes not be the best way to deal with the needs of individuals [43].

Several authors paint an even darker picture of the problem. They envision a smart city as a dystopian “city-as-a-machine” and an “omnipresent Truman Show.” They claim that many people think that companies that implement smart city projects see citizens as barriers, not partners [5].

As illustrated in [21], this lack of concern for the opinions of ordinary citizens could lead to a misalignment between their real needs and what government officials see as important. They present a survey-based importance ranking of various topics related to urban life, showing differences between the opinions of average citizens and the opinions of authorities. Starting from the most important things, the priorities of the citizens were (1) living, (2) mobility, (3) environment, (4) governance, and

(5) economy. In contrast, government officials ranked their priorities as follows: (1) living, (2) environment, (3) economy, (4) mobility, and (5) governance.

---

## 1.5 Models of Decision-Making and Citizen Engagement

Ma et al. [32] derived a simple model of how smart cities make data-driven decisions. The model is depicted in Fig. 1.3. It has three main components. First, a city collects data through various types of sensors. Second, the “smart services” process collected data. Third, “actuators” implement the decisions made by these smart services.

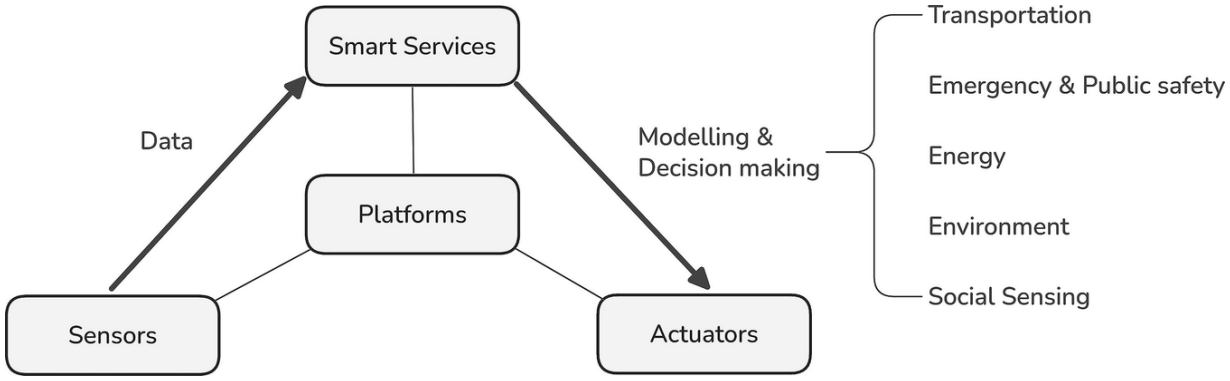


***Fig. 1.3*** A model of decision-making in smart cities [32]

Humans can assume any of the three roles in the model. This is consistent with Pierce et al. [36] who say that humans can act both as information producers and as information consumers. However, some authors see human involvement as a challenge. Humans bring “messiness” because their behavior is uncertain and unpredictable but, on the other hand, effective in computing with words without precise measurements [19].

Farinosi et al. [13] describe three levels of citizen participation.

1. Informative: cities inform their citizens about what is happening

2. Advisory: active listening and collection of opinions and ideas

3. Deliberative: involvement of the public in actual decision-making

Berntzen and Johannessen [7] distinguish between two roles citizens assume in the decision-making process:

- Agenda-setting
- Decision-making

When implementing participation policies and mechanisms, municipalities should think about the expected output and about the unique competences and experiences of individual citizens that municipalities can use in their favor.

French and Folz [15] explored how different types of municipal executives consult other stakeholder groups. They found that city managers consult other stakeholder groups more than mayors. The next observation is that the public is less involved than other groups, such as the city council or department heads.

The proposed smart cities maturity model [44] emphasizes that citizen eEmpowerment can be achieved through the support of their participation [45]. In this direction, the support of ICTs and smooth communication by soft computing between governments and citizens should be established. Citizens should not be mere absorbers of information but should be able to communicate and participate effectively. The framework consists of five levels: *eInforming*, *eConsulting*, *eDiscussion*, *eParticipation*, and *eEmpowerment* shown in Fig. 1.4. The comparative size of the city administration is

gradually reduced with moving into more sophisticated levels.

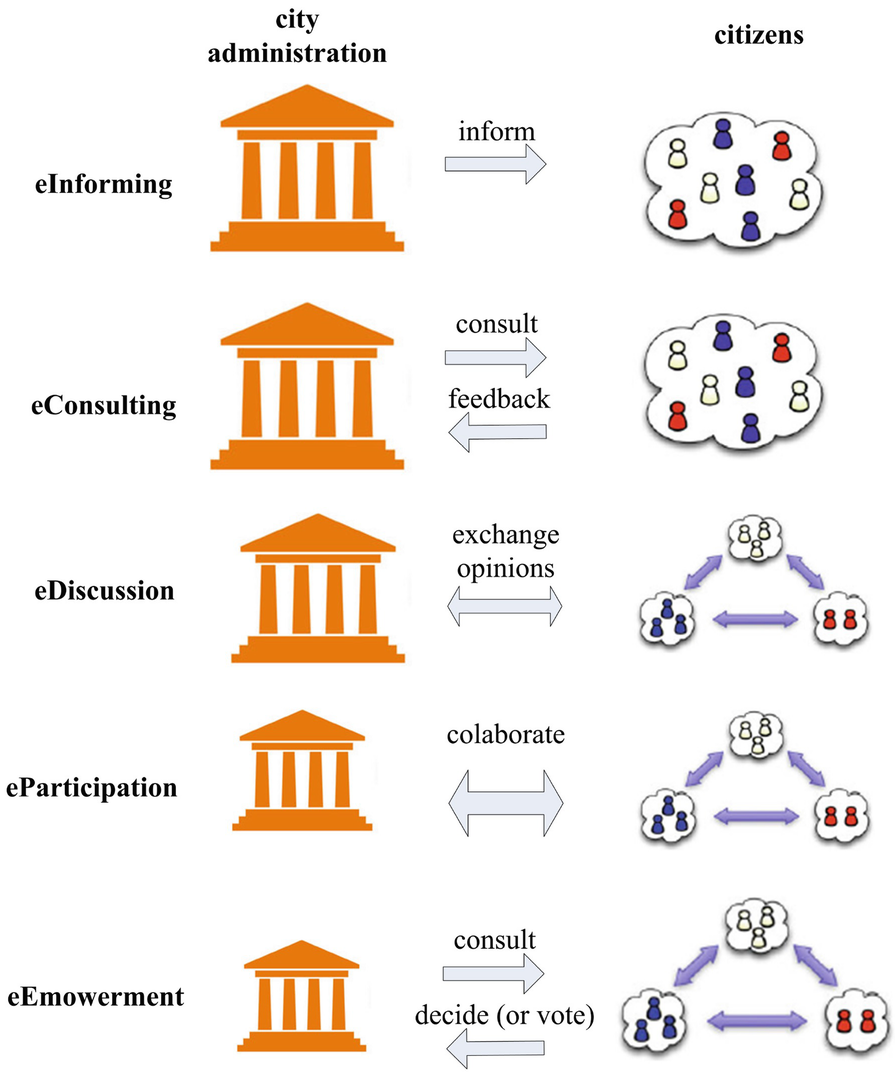


***Fig. 1.4*** Maturity model of smart cities adapted from [44, 45]

Smart cities should include eParticipation when developing new services for citizens and motivate citizens to use traditional and also virtual channels to express their opinions and suggestions. At the fifth level a bi-directional channel manages the decision in the hands of the citizens, that is, implementing what they have decided based on informing, consulting with city administration, and consulting among citizens.

The benefits of fuzzy logic are recognized and proposed at all levels. For example, at the first level effective communication can be handled by linguistic summaries, among other fuzzy logic approaches [25] to enhance explainability and interpretability of statistical figures.

The need for citizens to receive all relevant data in an understandable form, to freely express their opinions, and to have the right to participate in the decision-making process is a key feature for a prosperous city. Social networks have opened new channels for free discussion of all issues. It causes a shift away from the traditional ways, such as absorbing information from newspapers and sending posts or opinions to the editor for possible publication. However, posts on social networks are rapidly growing, so approaches based on computational intelligence should be applied.

Citizens, local government bodies, journalists, and other stakeholders face two main challenges in an intensive data-driven society [2]: (1) extracting valuable knowledge from diverse data sets and (2) becoming active actors, not mere absorbers. Solutions based on fuzzy sets and logic are significant in the context of smart cities due to their ability to robustly envelope semantic uncertainty, imprecise concepts, and vague data, which are a common sign in any human interaction.

Citizen engagement is necessary for developing effective governance strategies that ensure that new or modified acts and regulations address the needs of the citizens. The fact that citizens are not always able to clearly express their votes, like referendum *yes*, *no* type voting, should not be neglected. Thus, a flexible referendum voting model for collecting and interpreting results has been proposed in [24]. Consequently, an experiment is performed in Basel [39], revealing that the flexible referendum also attracts less decisive voters to provide their opinion as an

inclination toward clear acceptance or clear rejection. Referendum questions are created from different inputs; one of them might be the result of a deep evaluation of sentiments posted on social media.

In complex tasks, it is advisable to create a flexible transparent model to handle and evaluate citizen votes. For example, in a process of selecting the most suitable route for a new road within a municipality area, citizens cannot clearly express their support or rejection of each proposed alternative, and, moreover, the alternatives affect different subgroups of citizens [12]. The steps were as follows: Citizens were informed, they consulted municipal representatives, exchanged opinions, on their meetings they collaborated, and finally each citizen expressed its intensity of support and rejection. By three-level aggregation, the most suitable alternative is recognized. The relevant task is to motivate citizens to explore alternatives and improve usage level [6].

Participation in democratic processes is usually not high. For example, Swiss citizens in most cases participate less than 50% at the municipal level [38]. Many citizens find it challenging to express their votes or preferences in a rigid system where usually only absolute choices (*yes*, *no*) are allowed. Instead, citizens prefer to evaluate options using a spectrum of opinions, often relying on additional arguments, contextual information, and the existence of a way to flexibly explain the vote [12, 41]. However, implementing such systems requires careful design to ensure accessibility, transparency, intuitive use, and trust among voters (e.g., [11]).

On the other hand, citizens feel more free to express their opinions and react to the opinions of other citizens on social networks. The evaluation of opinion on social networks is a challenging task or *the journey from text to understanding*. Citizen posts are neither votes (classical

referendum of flexible referendum) nor scientific work abstracts. Posts might contain irony, short replies, long replies, grammar issues, double meaning, and the like. In addition, posts cannot simply be evaluated as positive or negative. The inclination or intensity of positiveness or negativeness is relevant. Few very negative posts are not the same as a bit higher number of lightly positive posts.

Instead of focusing on the level of participation, the study [7] distinguishes between two roles citizens might assume in the decision-making process: establishing the agenda and making decisions. When implementing participation policies and mechanisms, municipalities should think about expected output.

Two points of view in the use of Internet-based tools [7] are emphasized: Some see online tools as a replacement for traditional democratic processes, while others see them as a complement. Both perspectives are rather optimistic about the application of online technologies in the domain of municipal decision-making. For example, people who cannot attend various offline activities due to time and space constraints to express their opinions can use online platforms to express their opinions. However, people without access to the Internet prefer traditional approaches.

Many cities try to increase their “smartness” by developing smartphone applications. The study [3] explored how demographic factors affect the interest of citizens in using these applications. They discovered that less educated people are more likely to download and use smartphone applications than more educated people; however, they also note that the underlying cause could be the lower age. In terms of gender, women use municipal smartphone applications more than men. Next, three main groups of municipal application users are identified: adopters, worriers, and apathetic users. Local governments should not assume that people will readily use the

developed applications. The role of target marketing and the involvement of citizens in the development process should not be neglected.

The relationship between smart city technologies and people in general is discussed in [23]. These technologies have two main purposes: to help people in their personal lives and to help people participate in solving local issues. That work refers to several authors who claim that people are increasingly disconnected from local democratic structures. Civic engagement means both active participation and a sense of belonging. As such, it provides psychological benefits for both society and individuals. The engagement should be voluntary, which requires responsible institutions and mutual trust.

Finally, the work [10] considers the citizens as the most valuable element in a smart city. As such, they should be consulted on the issues the city is facing. To achieve this, a citizen engagement maturity model with three levels is proposed as follows:

1. **A disperse city:** Local authorities make most decisions.

2. **An integrated city:** Citizens' opinions are collected, traditionally or through some digital platforms.

3. **A connected city:** high degree of cooperation and coordination between local authorities and various citizen groups

In addition, the advantages of high civic participation include increased loyalty, a better reputation of the city, and greater attractiveness to skilled workers.

## 1.6 New Trends in Citizen Engagement

Recent literature points to several trends in the field of citizen participation. There are advances in both technologies and how the issue is perceived from the societal point of view.

**An Alternative View of a Smart City** A somewhat dire view on a smart city, a technological dystopia, is discussed in [5], where the authors contrast this perception with an alternative view according to which a truly smart city is based on participatory, community-based smart initiatives. Other authors such as [22] or [1] present a similar vision.

The alternative view emphasizes citizen engagement beyond a usual delivery of services by bottom-up democracy supported with modern technologies, greater social inclusion, power shift to ordinary citizens and their communities, identifying a problem first and finding a technological solution second, or respecting cultural and historical context. On the technological side, the alternative view finds support in open-source and free software initiatives. However, it is important to recognize that the movement is still relatively new and not widely accepted.

The work [5] equates the alternative view with crowdsourcing. When used in the context of a private enterprise, crowdsourcing usually refers to the act of taking a task usually performed by employees and outsourcing it to an unspecified network of people, usually in the form of an open call. The authors think that opinions and decision-making can be outsourced in the same way.

**Pervasive Computing** Pervasive computing describes the idea of embedding computational capabilities into everyday objects to minimize friction for users who interact

with computing devices. In the context of a smart city, this usually means devices that are available at the right place and time. Examples include displays at public transport stops or simple “thumbs-up or thumbs-down” feedback devices. Such pervasive technologies might play an important role in citizen engagement [23]. It is important for these devices to be interactive to facilitate *discussion in space*. For example, they could allow citizens to leave a message for the authorities or to express (dis)agreement with ideas and statements made by local politicians. In ancient times, Athenians left their opinion on clay plates and even wrote names of politicians who should leave their positions and city [17].

**Combining Multiple Means of Communication** Cities should not rely on any single technology and should instead support multiple means of communication, modern and traditional. Citizens are not always willing or capable to participate using the currently offered way of communication [23]. When it comes to using web-based technologies, cities can, for example, combine two different types of tools: generally available social media or proprietary platforms. Options for providing feedback offline should also be used. One challenge cities should address is that both online discussions and offline workshops tend to be dominated by activist groups, who often oppose planned projects. In contrast, local governments often fail to present their side of the story [7].

**Application of Sentiment Analysis and Topic Modeling Techniques** Word2Vec embeddings and an LDA topic modeling method to automatically categorize citizen complaints (leading to a reduction in the workload of municipal employees) are explored in [27] and briefly discussed the potential usefulness of dynamic topic modeling: monitoring changes in topic popularity over

time. A system for analyzing sentiment on Twitter is proposed in [5], an idea closely resembling the intended use of the proposed framework in this book. However, that work does not explicitly address the issue of opinion diversity, focusing only on a predetermined set of topics without examining multiple alternative approaches to sentiment analysis and topic modeling.

**Proper Use of Existing Technologies** Current computer-based opinion collection tools for simply replicating traditional methods were criticized in [33]. However, that work also shows how modern technologies could be used to improve these methods. A simple survey tool that lets respondents add custom answers that are then shown to other respondents was also presented. The idea is still in its infancy, and the authors themselves list many limitations. Most importantly, the method works well only if the number of respondents is small, while custom responses are limited in length. The trade-off between answer independence and helping citizens express themselves by having access to answers provided by other respondents is highlighted.

---

## 1.7 From Smart to Cognitive Cities

The concept of *cognitive cities* was developed in response to some of the aforementioned criticisms of smart cities. Two specific aspects are mentioned [43]: (1) the top-down approach to planning and (2) although the focus on efficiency and sustainability is important, it may not be the best way to deal with the needs of individuals. Citizens have a legitimate claim to participate in municipal decision-making. In this sense, the idea of cognitive cities covers and further extends many ideas from the previous section.

According to the literature review [31], the term *cognitive city* was first used by [35]. However, the first

term used in the context of architecture and urban planning was in [47]. It describes cognitive cities as *a "place" that enables involvement, generates awareness, uses creativity, deals with uncertainty, looks at centrality, and implies sustainability for future urban habitat questions*.

In the contemporary perspective, cognitive cities are based on the idea that cities need collective *urban intelligence* [43]. This collective intelligence consists of many *individual intelligences* that include not only human actors but also sensors or artificial intelligence systems. Collective intelligence emerges when individual intelligences communicate and interact with each other, with the human-in-the-loop approach being seen as crucial. Then a city acts and evolves as a superorganism [14].

The work [31] mentions *connectivism* as a theoretical substrate on which the idea of cognitive cities is based. Connectivism as a new theory of learning is introduced as [42]:

> Learning can reside outside of ourselves, is focused on connecting specialized information sets, and the connections that enable us to learn more are more important than our current state of knowing.

Some of the main principles of connectivism include the following:

- Learning and knowledge reside in diversity of opinions.
- Learning occurs by connecting specialized information sources.
- Learning may include nonhuman appliances.
- The maintenance of connections is crucial for continuous learning.
- Decision-making is a learning process.

An umbrella term *cognitive computing* is used in [26] to describe such systems. Cognitive computing is characterized by intelligence amplification loops, i.e., an enhancement of human cognitive capacities through computational tools. Among others, these computational tools include methods from the fields of machine learning and artificial intelligence. Soft computing methods also play an important role, as they are able to deal with uncertainty, human evaluation properties, and inferring without precise measurements. Special emphasis is on tools built on fuzzy set theory, such as linguistic summaries described in detail, for example, by [25].

According to [43], cognitive cities are not a replacement for smart cities. Instead, these two concepts should be considered complementary. They can coexist side by side and enhance each other's capabilities. In a sense, the concept of a cognitive city is just a different approach to smart city issues that emphasizes the role of interaction and communication. This vision is consistent with previously discussed work [5] and [1] who concluded that an ideal smart city combines a top-down approach driven by technology with a bottom-up approach driven by community.

Given this description, it can be claimed that the topic of this work is related to cognitive cities. The goal is to capture the knowledge emerging from a set of interconnected human actors. Moreover, to achieve this vision, the proposed framework utilizes various soft computing and machine learning methods.

---

# References


1. Abbott C (2020) City planning: a very short introduction. Oxford University Press, Oxford. https://doi.org/1093/actrade/9780190944346.001.0001 [Crossref]

2. Alonso JM, Castiello C, Mencar C (2019) The role of interpretable fuzzy systems in designing cognitive cities. In: Portmann E, Tabacchi M, Seising R, Habenstein A (eds) Designing cognitive cities. Springer, Cham, pp. 131–152
[Crossref]
3. Alcaide Muñoz L, Rodríguez Bolívar MP (2019) Demographic profile of citizens' interest, evaluation and opinions of local government apps in smart cities. In: Lindgren I, Janssen M, Lee H, Polini A, Rodríguez Bolívar M P, Scholl H J, Tambouris E (eds) Electronic government. Springer, Cham, pp 313–325. https://doi.org/1007/978-3-030-27325-5
[Crossref]
4. Andra F, Supangkat SH, Kosala R R (2018) Smart governance as smart city critical success factor (case in 15 cities in Indonesia). In: 2018 international conference on ICT for smart society (ICISS), Semarang, IEEE, pp 1–6. https://doi.org/10.1109/ICTSS.2018.8549923
5. Alizadeh T, Sarkar S, Burgoyne S (2019) Capturing citizen voice online: enabling smart participatory local government. Cities 95. 0264–2751. https://doi.org/10.1016/j.cities.2019.102400. https://www.sciencedirect.com/science/article/pii/S0264275118315981
6. Bach MP, Žmuk B, Stjepić AM, Khawaja S, Qureshi FH (2025) Citizen-centric insights into e-democracy and e-government: hybrid approach using cluster and multiple regression analysis. Central Eur J Oper Res. https://doi.org/1007/s10100-025-00963-2
7. Berntzen L, Johannessen MR (2016) The role of citizen participation in municipal smart city projects: lessons learned from Norway. In: Gil-Garcia JR, Pardo TA, Nam T (eds) Smarter as the new urban agenda: a comprehensive view of the 21st century city. Springer, Cham, pp 299–314.https://doi.org/10.1007/978-3-319-17620-8-16
[Crossref]
8. Bulukova D, Tabakovic M, Wahl H (2016) Smart cities education as mobility, energy & ICT hub. In: 2016 5th international conference on smart cities and green ICT systems (SMARTGREENS), pp 1–8
9. Capdevila I, Zarlenga M (2015) Smart city or smart citizens? The Barcelona case. J Strategy Manag 8:266–282. https://doi.org/2139/ssrn.2585682
10. Ceballos GR, Larios VM (2016) A model to promote citizen driven government in a smart city. In: 2016 IEEE international smart cities conference (ISC2), Trento, Italy. IEEE, pp 1–6. https://doi.org/1109/ISC2.2016.7580873
11. Colombo M (2024) Phenotropic interaction. Springer, Cham

[Crossref]
12. Emmenegger B, Hudec M, Zapletal F, Bigea G, Portmann E (2024) Application of the three-level aggregation model for evaluating opinions under hesitance for fuzzy voting in spatial planning public decision-making. In: 42nd international conference on mathematical methods in economics (MME 2024). The Czech Society for Operations Research, pp. 82–87
13. Farinosi M, Fortunati L, O'Sullivan J, Pagani L (2020) Capturing citizens' opinions through a combination of survey and online social data. Int J Commun 14:26
14. Finger M, Portmann E (2016) What are cognitive cities? In: Finger M, Portmann E (eds) Towards cognitive cities: advances in cognitive computing and its application to the governance of large urban systems. Springer, Cham, pp 1–11. https://doi.org/1007/978-3-319-33798-2_1
15. French PE, Folz D H (2004) Executive behavior and decision making in small U.S. cities. Am Rev Public Admin 34(1):52–66. https://doi.org/1177/0275074003259186
[Crossref]
16. Foltynova HB, Rybova K, Vejchodska E (2019) Reflection of smart city approaches in transport strategic planning in Czech cities. In: 2019 Smart city symposium Prague (SCSP), Prague, Czech Republic. IEEE, Piscataway, pp 1–6. https://doi.org/10.1109/SCSP.2019.8805714
17. Gannon M (2022) Ancient Greeks voted to kick politicians out of Athens if enough people didn't like them. Smithsonian Magazine, 13
18. Gobson DV, Kozmetsky G, Smilor WG (1992) The technopolis phenomenon: smart cities, fast systems, global networks
19. Giffinger R, Fertner C, Kramar H, Kalasek R, Milanović N, Meiers E (2007) Smart cities—ranking of European medium-sized cities
20. Greengard S (2023) Internet of Things. Encyclopædia Britannica. https://www.britannica.com/science/Internet-of-Things
21. Gupta K, Hall RP (2017) The Indian perspective of smart cities. In: 2017 smart city symposium prague (SCSP), Prague, Czech Republic. IEEE, Piscataway, pp 1–6. https://doi.org/10.1109/SCSP.2017.7973837
22. Halegoua G R (2020) Smart cities. The MIT Press, Cambridge
[Crossref]
23. Hosio SJ, Gonçalves J, Kostakos V, Riekki J (2015) Crowdsourcing public opinion using urban pervasive technologies: lessons from real-life experiments in Oulu. Policy & Internet 7:203–222. https://doi.org/1007/978-3-030-27325-5

[Crossref]
24. Hudec M, Mináriková E, Schwarz D, Fivaz J (2022). Estimating linguistic summaries on the unit interval data. In: Proceedings of the 2022 IEEE international conference on fuzzy systems (FUZZ-IEEE), Padua, Italy, pp. 1–8
25. Hudec M (2019) Possibilities for linguistic summaries in cognitive cities. In: Portmann E, Tabacchi ME, Seising R, Habenstein A (eds) Designing cognitive cities. Springer, Cham, pp 47–84
[Crossref]
26. Kaltenrieder P, Portmann E, Myrach T (2015) Fuzzy knowledge representation in cognitive cities. In: 2015 IEEE international conference on fuzzy systems (FUZZ-IEEE), Istanbul, Turkey. IEEE, Piscataway, pp 1–8
27. Kim B, Yoo M, Park KC, Lee KR, Kim JH (2021) A value of civic voices for smart city: a big data analysis of civic queries posed by Seoul citizens. Cities 108:102941. https://doi.org/1016/j.cities.2020.102941
[Crossref]
28. Kumar TMV (2015) E-governance for smart cities. Springer, Berlin, pp 1–43. https://doi.org/1007/978-981-287-287-6
29. Kumar, TMV (eds) (2017). Smart economy in smart cities: International collaborative research: Ottawa, St. Louis, Stuttgart, Bologna, Cape Town, Nairobi, Dakar, Lagos, New Delhi, Varanasi, Vijayawada, Kozhikode, Hong Kong. Springer, Singapore
30. Koutra S, Becue V, Ioakimidis CS (2018) A multiscalar approach for “smart city” planning. In: 2018 IEEE international smart cities conference (ISC2), Kansas City. IEEE, Piscataway, pp 1–7. https://doi.org/10.1109/ISC2.2018.8656889
31. Machin J, Solanas A (2018) A review on the meaning of cognitive cities. In: 2018 9th international conference on information, intelligence, systems and applications (IISA), Zakynthos. IEEE, Piscataway, pp 1–5
32. Ma M, Preum SM, Ahmed MY, Tarneberg W, Hendawi AM, Stankovic JA (2019) Data sets, modeling, and decision making in smart cities. ACM Trans Cyber-Phys Syst 4:1–28. https://api.semanticscholar.org/CorpusID:204897889
[Crossref]
33. Matsubara M, Matsuda Y, Kuzumi R, Koizumi M, Morishima A (2020) Collecting and organizing citizen opinions: a dynamic microtask approach and its evaluation. In: Sundqvist A, Berget G, Nolin J, Skjerdingstad K I (eds) Sustainable digital communities. Springer, Cham, pp 50–65
[Crossref]

34. Nieves JP (2022) Smart urban logistics: Improving delivery services by computational intelligence. Springer, Cham
[Crossref]
35. Novak M (1997) Cognitive cities: intelligence, environment and space. In: Droege P (ed) Intelligent environments. North-Holland, Amsterdam, pp 386–420. https://doi.org/1016/B978-044482332-8/50023-8
[Crossref]
36. Pierce P, Andersson B (2017) Challenges with smart cities initiatives—a municipal decision makers' perspective. In: Hawaii international conference on system sciences. https://api.semanticscholar.org/CorpusID:1117249
37. Rifkin J (2016) How the third industrial revolution will create a green economy. New Perspect Q 33(1):6–10. https://doi.org/1111/npqu.12017
[Crossref]
38. Sciarini P, Tresch A (2023) Direct-democratic votes. In: Emmenegger P, FossatiF, Häusermann S, Papadopoulos Y, Sciarini P, Vatter A (eds) The Oxford handbook of Swiss politics. Oxford University Press, Oxford, pp. 410–429
39. Schwarz D, Emmenegger B, Fivaz J, Hudec M, Portmann E (2024) Effects of fuzzy voting on citizen participation: empirical evidence from Switzerland. In Proceedings of the tenth international conference on eDemocracy & eGovernment (ICEDEG), Lucerne, pp. 1–8
40. Shishanina MA, Sidorov AA (2021) Management of social and economic development of municipalities. arXiv preprint, arXiv:2111.13690
41. Schwarz D, Emmenegger B, Fivaz J, Hudec M, Portmann E (2024) Effects of fuzzy voting on citizen participation: empirical evidence from Switzerland. In: Tenth international conference on eDemocracy & eGovernment (ICEDEG 2024). IEEE, Piscataway, pp. 1–8
42. Siemens G (2004) Connectivism: a learning theory for the digital age. Int J Instruct Technol Distance Learn 2. https://api.semanticscholar.org/CorpusID:10302016
43. Tabacchi ME, Portmann E, Seising R, Habenstein A (2019) In: Portmann E, Tabacchi M E, Seising R, Habenstein A (eds) Designing cognitive cities. Springer, Cham, pp 3–27. https://doi.org/10.1007/978-3-030-00317-3_1
[Crossref]
44. Terán L, Kaskina A, Meier A (2016) Maturity model for cognitive cities. In: Portmann E, Finger M (eds) Towards cognitive cities. Springer, Cham

45. Terán L. (2014) SmartParticipation: a fuzzy-based recommender system for political community-building. In: Fuzzy management methods. Springer, Cham
46. Thorne C, Griffiths (2014) Smart, smarter, smartest: redefining our cities. In: Dameri RP, Rosenthal-Sabroux C (eds) Smart city. Springer, Cham, pp 89–99. https://doi.org/1007/978-3-319-06160-3_4 [Crossref]
47. Tusnovics DA (2007) Cognitive cities: interdisciplinary approach reconsidering the process of (re)inventing urban habitat. https://api.semanticscholar.org/CorpusID:201047107
48. Washburn D, Sindhu U, Balaouras S, Dines, RA, Hayes N, Nelson LE (2009) Helping CIOs understand “smart city” initiatives. Growth 17(2):1–17
49. Willis KS (2019) Whose right to the smart city? In: Cardullo P, Di Feliciantonio C, Kitchin R (eds) The right to the smart city. Emerald Publishing Limited, pp 27–41. https://doi.org/10.1108/978-1-78769-139-120191002
50. Yang J, Kwon Y, Kim D (2021) Regional smart city development focus: the South Korean National. IEEE Access 9:7193–7210. https://doi.org/10.1109/ACCESS.2020.3047139 [Crossref]

# Part I
# Methodology

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2025
M. Švaňa et al., *From Text to Understanding*, Fuzzy Management Methods
https://doi.org/10.1007/978-3-032-00129-0_2

# 2. Fuzzy Sets and Fuzzy Algebra

Miloš Švaňa[1] ✉, František Zapletal[1] and Miroslav Hudec[1]
(1) VSB – Technical University of Ostrava, Ostrava-Poruba, Czech Republic

## 2.1 Introduction to Fuzzy Mathematics

The first work on the fuzzy approach dates back to 1965 when L. A. Zadeh introduced the concept of fuzzy sets [25]. Today, it is hard to find the field of modeling under semantic uncertainty where the fuzzy approach has not been used so far. The motivation for using fuzzy sets in modeling is that they handle the uncertainty of various types and make models more realistic.

Many real-world problems, including decision-making based on natural language data, are under the impact of uncertainty. If it were possible to adequately capture and describe uncertainty instead of avoiding or ignoring it, we could come up with much better solutions to those problems. Several different sources of uncertainty can be distinguished:

- Even if something (event/set/element of reality) is measurable, it is impossible to provide perfectly precise value (temperature of the air outside).
- Sometimes the precise value is not possible because its high dynamics in time (actual demand on market).
- The precise value is not known because it will be determined in the future (price of energy tomorrow).
- Situations when it is necessary to save time and “sacrifice” the precise value (number of people skiing at a hill)
- The precise values can be measured but are not available (marketing department explores how many young people visit the shop—they cannot ask about the age of customers).
- The inability to define precise boundaries, especially when using the qualitative description (hot water, young man, and quite satisfied customer)
- Uncertainty in relations (I enjoy this car much more than another one)
- Hesitance in opinion (I do not know if I enjoyed the movie or not)

The sources of uncertainty listed above can also be combined. It is worth mentioning that the uncertainty can never be fully eliminated; on the other hand, not all uncertainty factors have to be necessarily captured by models. It is up to the skills and knowledge of a modeler and also to the purpose of modeling what will be captured in the description and what will not.

In the classical "crisp" set theory, the belonging of an element $x$ to the set is described by the characteristic function $\chi_A$:

$$\chi_A(x) : X \to \{0; 1\}, \tag{2.1}$$

where $X$ is a universal set and $x$ its element.

The core idea of the fuzzy set is that an element $x$ can belong to the fuzzy set not only fully but also partially, i.e., to some extent usually measured in the unit interval.

Fuzzy logic is an extension of infinite-valued logic by incorporating fuzzy sets into the system of infinite-valued logic [6]. Fuzzy logic is related to $T$-norms as conjunction, to $T$-conorms as disjunction, and to the unit interval as a set of truth values. The principle of extensionality says that the truth value of a compound sentence depends only on the truth values of the elementary components and extends crisp mathematical functions to fuzzy sets allowing, for example, operations like addition and multiplication of fuzzy numbers [9]. It is a valuable observation for diverse real-world tasks where elementary conditions might be of variety forms like binary condition, fuzzy concept and crisp data, crisp concept and fuzzy data, fuzzy concept and fuzzy data, and the like.

### 2.1.1 Fuzzy Sets, Their Description, and Properties

In order to maintain the clarity of the written text, all fuzzy sets will be denoted by tilde in this work. In line with [20], the following two definitions describe a fuzzy set and a membership function, respectively.

**Definition 2.1** Let $X$ be a non-empty set. A fuzzy subset $\tilde{A}$ of $X$ is the family of subsets $A_\alpha \subseteq X$ satisfying $A_0 = X, A_\beta \subseteq A_\alpha$, and $A_\beta = \bigcap_{\alpha \in [0,1]} A_\alpha$ for all $0 \leq \alpha < \beta \leq 1$. A fuzzy subset $\tilde{A} = \{A_\alpha\}_{\alpha \in [0,1]}$ of $X$ is called a *fuzzy set*.

**Definition 2.2** Let $\tilde{A} = \{A_\alpha\}_{\alpha \in [0,1]}$ be a fuzzy subset of $X$. The function $\mu_{\tilde{A}} : X \to [0, 1]$ defined by

$$\mu_{\tilde{A}}(x) = \sup \{\alpha \mid \alpha \in [0, 1], x \in A_\alpha\}$$

is called the *membership function* of $A$, and its value $\alpha$ at $x \in X$ is called the *membership degree* of $x$ in $A$.

The reader can see that the membership function really extends the concept of characteristic functions used for crisp sets. For further work

with fuzzy sets, it is useful to define aforementioned $\alpha$-cut: core and support of a fuzzy set.

**Definition 2.3** The *$\alpha$-cut* $[A]_\alpha$ of the fuzzy set $\tilde{A}$ is the upper-level set of its membership function $\mu_{\tilde{A}}(x)$ at membership level $\alpha \in [0,1]$:

$$[\tilde{A}]_\alpha = \{x \in X \mid \mu_{\tilde{A}}(x) \geq \alpha\}.$$

Following Definition 2.1 and Definition 2.3, a fuzzy set can be determined either by its membership function or as a collection of its $\alpha$-cuts for $\alpha \in [0,1]$.

**Definition 2.4** The *core* of $\tilde{A}$ is defined by

$$Core(\tilde{A}) = \{x \in X \mid \mu_{\tilde{A}}(x) = 1\}.$$

In other words, the core of a fuzzy set is the set of values for which it is claimed that they belong clearly to the set.

**Definition 2.5** The *support* of $\tilde{A}$ is defined by

$$Supp(\tilde{A}) = \{x \in X \mid \mu_{\tilde{A}}(x) > 0\}.$$

This means that the support covers all values for which a decision-maker at least admits that they can belong to the set.

**Definition 2.6** A fuzzy set $\tilde{A} = \{A_\alpha\}_{\alpha \in [0,1]}$ is called a *fuzzy interval* defined on $X$ (for the sake of simplicity, let $X \subseteq \mathbb{R}$) if its $\alpha$-cuts are non-empty convex sets for all $\alpha \in [0,1]$.

Due to the fact that all $\alpha$-cuts of a fuzzy interval $\tilde{A}$ are convex sets, they are described by compact intervals with lower bounds $[\tilde{A}]_\alpha^L$ and upper bounds $[\tilde{A}]_\alpha^R$. A crisp interval is a special case of the fuzzy interval with the core equal to its support. The most common type of fuzzy interval is then a trapezoidal interval $\tilde{A} = (a, m, n, b)$, which has a piecewise linear membership function (2.2)—formalizing concepts like *medium satisfaction*, see Fig. 2.1.

$$\mu_{\tilde{A}}(x) = \begin{cases} 0, & x < a \\ \frac{x-m}{m-a}, & a < x \leq m \\ 1, & m < x < n \\ \frac{b-x}{b-n}, & n \leq x < b \\ 0, & x > b. \end{cases} \tag{2.2}$$

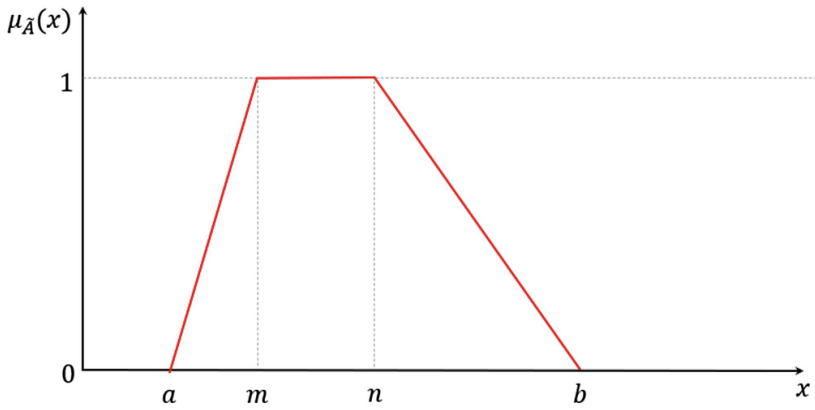


***Fig. 2.1*** Example of a trapezoidal fuzzy interval

**Definition 2.7** A fuzzy interval $\tilde{A} = \{A_\alpha\}_{\alpha \in [0,1]}$ is called a *fuzzy number* if its core is a singleton, i.e., $\mu_{\tilde{A}}(x) = 1$ for precisely one element of $x \in X$ .

In applications of fuzzy sets, it is common to work often with the simplest fuzzy numbers, that is, triangular fuzzy numbers (TFNs) $\tilde{A} = (a, m, b)$—to express observations such as the *size of the living room around 25* $m^2$, see Fig. 2.2:

$$\mu_{\tilde{A}}(x) = \begin{cases} \frac{x-a}{m-a}, & a \leq x \leq m \\ \frac{b-x}{b-m}, & m \leq x \leq b \\ 0, & \text{otherwise.} \end{cases} \tag{2.3}$$

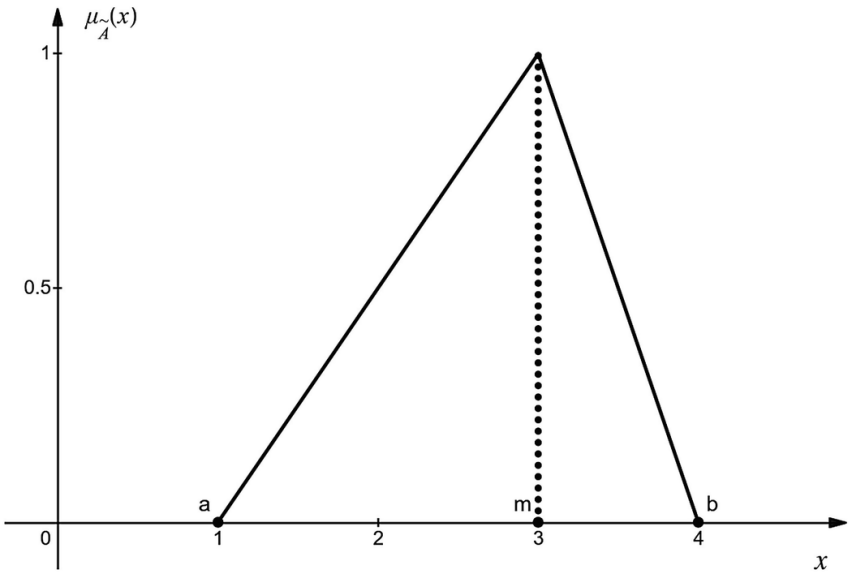


***Fig. 2.2*** Example of a triangular fuzzy number

A crisp (e.g., real) number is a special case of the fuzzy number with the core equal to its support.

**Definition 2.8** $hgt(\tilde{A}) = \max_{\alpha} \mu_{\tilde{A}}$ is called the *height* of the fuzzy set.

The height of any fuzzy interval must be equal to 1.

---

## 2.2 Fundamentals of Fuzzy Algebra

When one has a given part of reality described by fuzzy sets, it is potentially very useful, but without any further processing and calculations, the benefits are not fully used. Some linguistical summarization can be required (e.g., *most of the young people have rather positive opinion about new city rule*), some comparisons of fuzzy sets are needed (the opinion about social networks is *much more positive than the opinion about streaming platforms*), or/and some numerical operations with fuzzy sets are needed too (e.g., the level of uncertainty is reduced when the information is confirmed by expert).

Basic operations on sets like intersection or union, the relation of comparison (less, equal, or greater), or basic arithmetical operations with numbers (addition, multiplication) are trivial when working with crisp sets and numbers, respectively. However, when switching to the fuzzy environment, their fuzzy extensions must be used instead. These extensions are based on elementary operations between fuzzy sets called triangular norms ($T$-norms) for conjunction and triangular conorms ($T$-conorms) for disjunction. Next, the defuzzification measure is relevant when transforming the fuzzy set into the crisp representative value (e.g., maxima methods [18], Choquet integral [11], or gravity methods

considering the entire shape of the membership function [23]). However, defuzzification can lead to the loss of a portion of valuable information. Therefore, it should be decided whether to keep uncertainty in the whole process or to defuzzify into crisp counterparts in a particular computational stage. For more detailed information, see [19].

At first, the basic operations of fuzzy sets are introduced: intersection by $T$-norms and union by $T$-conorms [3, 10, 13].

**Definition 2.9** A class of functions $T : [0,1]^2 \to [0,1]$ that are commutative, associative, and monotonic and with the neutral element 1 is called the $T$-norms.

The properties of the $T$-norms can be written as

$$T(a,1) = a \quad \text{(boundary condition)}, \tag{2.4}$$
$$b \le c \Rightarrow T(a,b) \le i(a,c) \quad \text{(monotonicity)}, \tag{2.5}$$
$$T(a,b) = T(b,a) \quad \text{(commutativity)}, \tag{2.6}$$
$$T(a,T(b,c)) = T(T(a,b),c) \quad \text{(associativity)}. \tag{2.7}$$

**Definition 2.10** A class of functions $S : [0,1]^2 \to [0,1]$ that are commutative, associative, and monotonic with the neutral element 0 is called the $T$-conorms.

The properties of the $T$-conorms can be written as

$$S(a,0) = a \quad \text{(boundary condition)}, \tag{2.8}$$
$$b \le c \Rightarrow S(a,b) \le i(a,c) \quad \text{(monotonicity)}, \tag{2.9}$$
$$S(a,b) = S(b,a) \quad \text{(commutativity)}, \tag{2.10}$$
$$S(a,S(b,c)) = S(S(a,b),c) \quad \text{(associativity)}. \tag{2.11}$$

It can be seen that the only difference in properties between the $T$-norms and the $T$-conorms is the boundary condition, i.e., neutral elements 0 and 1. An interested reader can also check that the classical "crisp" intersection also meets the properties of the $T$-norms, and the classical "crisp" union also meets the properties of the $T$-conorms. However, unlike the crisp set theory, there are many functions $T$ and $S$ applied to fuzzy sets. In other words, more different $T$-norms and $T$-conorms exist and are applied. Their suitability mostly depends on the opinion of a decision-maker and further required properties. The representative $T$-norms (idempotent, strict, nilpotent, and drastic) are

$$\min \text{ (Gödel's) } T\text{-norm:} \quad T(a,b) = \min\{a,b\}, \tag{2.12}$$
$$\text{Product } T\text{-norm:} \quad T(a,b) = ab, \tag{2.13}$$
$$\text{Lukasiewicz } T\text{-norm:} \quad T(a,b) = \max(0, a+b-1), \tag{2.14}$$

$$\text{Drastic } T\text{-norm:} \quad T(a,b) = \begin{cases} b & \text{if } a = 1, \\ a & \text{if } b = 1, \\ 0 & \text{otherwise.} \end{cases} \tag{2.15}$$

The representative $T$-conorms, dual to t-norms, are

$$\text{max (Gödel's) } T\text{-conorm:} \quad S(a,b) = \max\{a,b\}, \tag{2.16}$$

$$\text{Product } T\text{-conorm:} \quad S(a,b) = a + b - ab, \tag{2.17}$$

$$\text{Lukasiewicz } T\text{-conorm:} \quad S(a,b) = \min(1, a+b), \tag{2.18}$$

$$\text{Drastic}t\text{-conorm:} \quad S(a,b) = \begin{cases} b & \text{if } a = 0, \\ a & \text{if } b = 0, \\ 1 & \text{otherwise.} \end{cases} \tag{2.19}$$

For other $T$-norms and $T$-conorms, one can look in, e.g., [3, 10, 12, 13].

In this text, only Gödel's $T$-(co)norm is considered to get the intersection and union of two fuzzy sets $\tilde{A}$ and $\widetilde{B}$:

$$\mu_{\tilde{A}\cap\widetilde{B}}(x) = \min\{\mu_{\tilde{A}}(x), \mu_{\widetilde{B}}(x)\} \tag{2.20}$$

and standard union as

$$\mu_{\tilde{A}\cup\widetilde{B}}(x) = \max\{\mu_{\tilde{A}}(x), \mu_{\widetilde{B}}(x)\}. \tag{2.21}$$

To complete the algebra for fuzzy sets, the complement operation is required. The standard complement, defined by Zadeh [25] is as follows:

$$\mu_{\tilde{A}'}(x) = 1 - \mu_{\tilde{A}}(x). \tag{2.22}$$

The operations defined previously are illustrated in Fig. 2.3.

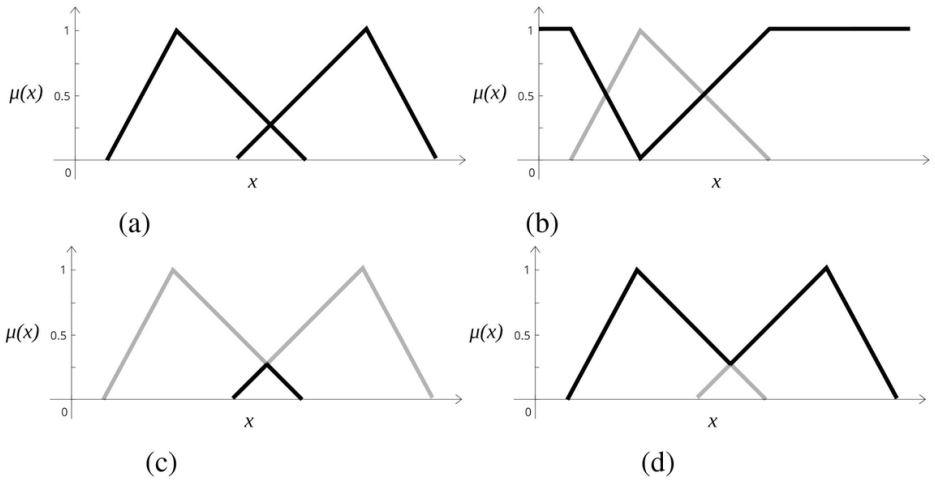


***Fig. 2.3*** Example of a standard complement, Gödel's intersection, and Gödel's union. (**a**) Two original sets. (**b**) Complement of the left set. (**c**) Intersection. (**d**) Union

Now, it is possible to proceed to fuzzy extensions of other crisp functions. In line with [19], the fuzzy extensions of crisp relations can be defined using the so-called Zadeh extension principle, based on $T$-norms and $T$-conorms, generally defined as follows:

**Definition 2.11** Let $T$ be a $T$-norm and $S$ be a $T$-conorm. Then, let $R$ be a crisp binary-valued relation on $X$. Fuzzy extensions $R^T$ and $R^S$ of the relation $R$ on $X$ defined for all fuzzy sets $\tilde{A}$ and $\widetilde{B}$ with the membership functions $\mu_{\tilde{A}} : X \to [0,1]$ and $\mu_{\widetilde{B}} : X \to [0,1]$, respectively, by

$$\mu_{R^T}(\tilde{A},\widetilde{B}) = \sup\left\{T\Big(T\big(\mu_{\tilde{A}}(x),\mu_{\widetilde{B}}(y)\big),\mu_R(x,y)\Big) \mid x,y \in X\right\}, \tag{2.23}$$

$$\mu_{R^S}(\tilde{A},\widetilde{B}) = \inf\left\{S\Big(S\big(1-\mu_{\tilde{A}}(x),1-\mu_{\widetilde{B}}(y)\big),\mu_R(x,y)\Big) \mid x,y \in X\right\} \tag{2.24}$$

are called a $T$-fuzzy extension of the crisp relation $R$ and $S$-fuzzy extension of the crisp relation $R$, respectively.

As mentioned above, only Gödel's $T$-(co)norm will be further used in this book, see (2.20) and (2.21). Thus, the fuzzy extension can be written as follows:

$$\mu_{R^{\min}}(\tilde{A},\widetilde{B}) = \sup_{R(x_1,x_2)}\left\{\min\left(\mu_{\tilde{A}}(x_1),\mu_{\widetilde{B}}(x_2)\right) \mid x_1,x_2 \in \mathbb{R}\right\}, \tag{2.25}$$

$$\mu_{R^{\max}}(\tilde{A},\widetilde{B}) = \inf_{R(x_1,x_2)}\left\{\max\left(1-\mu_{\tilde{A}}(x_1),1-\mu_{\widetilde{B}}(x_2)\right) \mid x_1,x_2 \in \mathbb{R}\right\}. \tag{2.26}$$

Fuzzy extension of crisp addition to addition of fuzzy numbers is written as [19]

$$\tilde{A} \oplus \widetilde{B} = \mu_{\oplus(\tilde{A},\widetilde{B})}(z) = \sup_{z=x_1+x_2}\left\{\min\left(\mu_{\tilde{A}}(x_1),\mu_{\widetilde{B}}(x_2)\right) \mid x_1,x_2 \in \mathbb{R}\right\}, \tag{2.27}$$

while multiplication of a fuzzy number by a crisp number is [19]

$$k\tilde{A} = \mu_{k\cdot\tilde{A}}(x) = \left\{\mu_{\tilde{A}}(k\cdot x) \mid x \in X, k \in \mathbb{R}\right\}. \tag{2.28}$$

**Remark 2.1** If $\tilde{A}$ and $\widetilde{B}$ in (2.27) and (2.28) are triangular fuzzy numbers $\tilde{A}(c^{\tilde{A}}-\sigma_L^{\tilde{A}},c^{\tilde{A}},c^{\tilde{A}}+\sigma_R^{\tilde{A}})$ and $\widetilde{B}(c^{\widetilde{B}}-\sigma_L^{\widetilde{B}},c^{\widetilde{B}},c^{\widetilde{B}}+\sigma_R^{\widetilde{B}})$, then it simplifies fuzzy extension [5] of addition as

$$\tilde{A} \oplus \widetilde{B} = (c^{\tilde{A}}+c^{\widetilde{B}}-\sigma_L^{\tilde{A}}-\sigma_L^{\widetilde{B}},c^{\tilde{A}}+c^{\widetilde{B}},c^{\tilde{A}}+c^{\widetilde{B}}+\sigma_R^{\tilde{A}}+\sigma_R^{\widetilde{B}})$$

and multiplication by a crisp parameter as

$$k\tilde{A} = (k\cdot(c^{\tilde{A}}-\sigma_L^{\tilde{A}}),k\cdot c^{\tilde{A}},k\cdot(c^{\tilde{A}}+\sigma_R^{\tilde{A}})), \qquad \text{for } k>0.$$

Several possibilities for comparing fuzzy numbers are presented in, e.g., [5, 19, 20]. A natural way to define a fuzzy ordering in fuzzy sets is realized by minimum and maximum fuzzy extensions of crisp binary relations $\leq$ and $\geq$, respectively. Based on (2.25) and (2.26), relations $\tilde{A} \lesssim^{\min} \widetilde{B}$ and $\tilde{A} \lesssim^{\max} \widetilde{B}$ can be defined as follows:

$$\mu_{\lesssim^{\min}}(\tilde{A},\widetilde{B}) = \sup_{x_1\leq x_2}\left\{\min\left(\mu_{\tilde{A}}(x_1),\mu_{\widetilde{B}}(x_2)\right) \mid x_1,x_2 \in \mathbb{R}\right\}, \tag{2.29}$$

$$\mu_{\tilde{\leq}^{\max}}(\tilde{A},\widetilde{B}) = \inf_{x_1 \leq x_2} \left\{ \max\left(1-\mu_{\tilde{A}}(x_1), 1-\mu_{\widetilde{B}}(x_2)\right) \mid x_1, x_2 \in \mathbb{R} \right\}. \quad (2.30)$$

$\alpha$-cuts (Definition 2.3) can be used to rewrite (2.29) and (2.30) in the form of (2.31) and (2.32), respectively [20].

$$\mu_{\tilde{\leq}^{\min}}(\tilde{A},\widetilde{B}) \geq \alpha \text{ if and only if } \inf [\tilde{A}]_\alpha \leq \sup [\widetilde{B}]_\alpha, \text{ thus if } [\tilde{A}]_\alpha^L \leq [\widetilde{B}]_\alpha^R; \quad (2.31)$$

$$\mu_{\tilde{\leq}^{\max}}(\tilde{A},\widetilde{B}) \geq \alpha \text{ if } \sup [\tilde{A}]_{1-\alpha} \leq \inf [\widetilde{B}]_{1-\alpha}, \text{ thus if } [\tilde{A}]_{1-\alpha}^R \leq [\widetilde{B}]_{1-\alpha}^L. \quad (2.32)$$

In order to make computations easier and more intuitive for decision-makers, the work [26] has introduced new measures corresponding to fuzzy extensions of the relational operators in (2.31) and (2.32), respectively. In particular, possibility and necessity measures have been established. The possibility measure corresponds to (2.31) and therefore can be interpreted as, for example, the possibility that the fuzzy set $\tilde{A}$ is less than or equal to the fuzzy set $\widetilde{B}$:

$$\mathrm{Pos}(\tilde{A} \leq \widetilde{B}) = \tilde{A} \,\tilde{\leq}^{\min}\, \widetilde{B} = \mu_{\tilde{\leq}^{\min}}(\tilde{A},\widetilde{B}). \quad (2.33)$$

The stricter fuzzy extension of the comparison relation $\leq$ given by (2.32) can be interpreted as a necessity measure, i.e., to what extent it is sure that the relation

$$\mathrm{Nec}(\tilde{A} \leq \widetilde{B}) = \tilde{A} \,\tilde{\leq}^{\max}\, \widetilde{B} = \mu_{\tilde{\leq}^{\max}}(\tilde{A},\widetilde{B}) \quad (2.34)$$

holds.

A graphical demonstration of a comparison of a fuzzy number with a crisp value and comparison of two fuzzy numbers using the possibility and necessity measures is provided by Figs. 2.4 and 2.5, respectively. For more fuzzy comparison relations, see, e.g., [8].

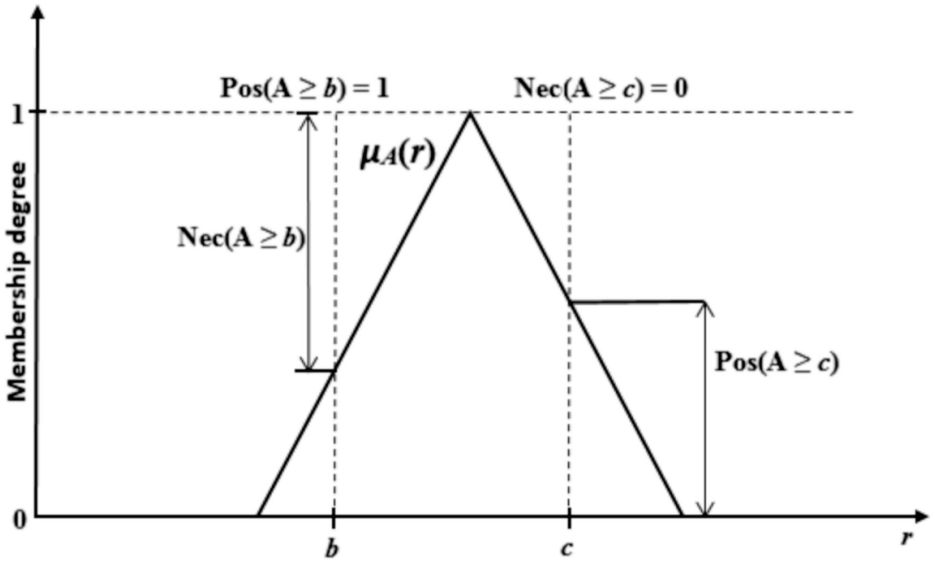


***Fig. 2.4*** Comparison of a fuzzy number $A$ and a crisp numbers $b, c$ using possibility and necessity measures

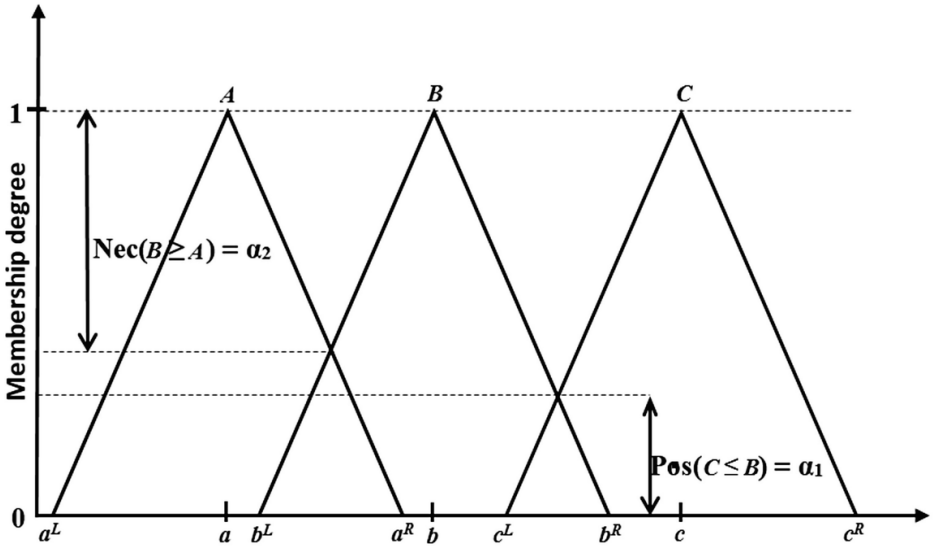


***Fig. 2.5*** Comparison of fuzzy numbers $A, B, C$ using possibility and necessity measures

From the perspective of sentiment analysis, the required level of adoption of fuzzy set theory and fuzzy algebra depends on the way the model is constructed and the way the data are processed. If the sentiment of individuals is expressed by TFN, it might be reasonable to reduce uncertainty by defuzzification and calculate the arithmetic mean. However, if one wants to handle uncertainty, then a fuzzy extension is applied. Next, by keeping the TFNs, the sentiment of two individuals can be mutually compared.

### 2.2.1 Similarity of Fuzzy Sets

In modeling, it is often necessary to know not only if (and to what extent) one set is greater/less than another but also how similar they are. For this purpose, several similarity measures, either based on distance measures between two fuzzy sets or derived from the fuzzy set theory concept, can be used.

The determination of a degree of similarity between two fuzzy sets is a crucial part of the data analysis framework proposed in this book. This section provides a brief overview of approaches to this problem.

A similarity measure $S(\tilde{A}, \widetilde{B})$ should have the following properties [4]:

- $0 \leq (\tilde{A}, \widetilde{B}) \leq 1$.
- $S(\tilde{A}, \widetilde{B}) = S(\widetilde{B}, \tilde{A})$.
- $S(\tilde{A}, \widetilde{A\prime}) = 0$, where $\widetilde{A\prime}$ denotes the complement of $\tilde{A}$.
- $S(\tilde{A}, \tilde{A}) = 1$.
- If $\tilde{A} \subseteq \widetilde{B} \subseteq \widetilde{C}$, then $S(\tilde{A}, \widetilde{B}) \geq S(\tilde{A}, \widetilde{C})$ and $S(\widetilde{B}, \widetilde{C}) \geq S(\tilde{A}, \widetilde{C})$.

In general, there are two main groups of fuzzy set similarity measures: distance-based and set-theoretic [4].

### *Distance-Based Similarity Measures*

Distance-based similarity measures use well-established distance measures [16, 24]:

- Hamming distance

$$d_H(\tilde{A}, \widetilde{B}) = \sum_{i=1}^{n} |\mu_{\tilde{A}}(x_i) - \mu_{\widetilde{B}}(x_i)|. \tag{2.35}$$

- Normalized Hamming distance

$$d_{nH}(\tilde{A}, \widetilde{B}) = \frac{1}{n} \sum_{i=1}^{n} |\mu_{\tilde{A}}(x_i) - \mu_{\widetilde{B}}(x_i)|. \tag{2.36}$$

- Euclidean distance

$$d_E(\tilde{A}, \widetilde{B}) = \sqrt{\sum_{i=1}^{n} (\mu_{\tilde{A}}(x_i) - \mu_{\widetilde{B}}(x_i))^2}. \tag{2.37}$$

- Generalized Minkowski distance which also covers the Hamming distance when $r = 1$ and Euclidean when $r = 2$

$$d_r(\tilde{A}, \widetilde{B}) = [\sum_{i=1}^{n} |\mu_{\tilde{A}}(x_i) - \mu_{\widetilde{B}}(x_i)|^r]^{1/r}. \tag{2.38}$$

- Supremum distance

$$d_\infty(\tilde{A}, \widetilde{B}) = \sup_i |\mu_{\tilde{A}}(x_i) - \mu_{\widetilde{B}}(x_i)|. \tag{2.39}$$

These distance metrics are used to calculate different similarity measures [14]:

- Koczy similarity measure

$$S(\tilde{A}, \widetilde{B}) = \frac{1}{1+d(\tilde{A}, \widetilde{B})}. \tag{2.40}$$

- Williams and Steele similarity measure with a steepness parameter $\alpha$

$$S(\tilde{A}, \widetilde{B}) = e^{-\alpha d(\tilde{A}, \widetilde{B})}. \tag{2.41}$$

- Santini similarity measure

$$S(\tilde{A}, \widetilde{B}) = 1 - d(\tilde{A}, \widetilde{B}). \tag{2.42}$$

### *Set Theory-Based Similarity Measures*

Instead of relying on distance metrics, set theory-based similarity measures are based on concepts from set theory, such as overlap or symmetric difference. Examples include [27]:

- Restle's similarity

$$S(\tilde{A}, \widetilde{B}) = 1 - |\tilde{A} \square \widetilde{B}|, \tag{2.43}$$

where, if standard intersection and union are used,

$$\tilde{A} \square \widetilde{B}(x) = \max [\min (\mu_{\tilde{A}}(x), 1 - \mu_{\widetilde{B}}(x)), \min (\mu_{\widetilde{B}}(x), 1 - \mu_{\tilde{A}}(x))]. \tag{2.44}$$

- Gregson's similarity

$$S(\tilde{A}, \widetilde{B}) = \frac{|\tilde{A} \cap \widetilde{B}|}{|\tilde{A} \cup \widetilde{B}|} \tag{2.45}$$

- Enta's similarity

$$S(\tilde{A}, \widetilde{B}) = \sup_{x \in X} (\tilde{A} \cap \widetilde{B})(x). \tag{2.46}$$

Given several desirable properties explained in the later chapters, a measure proposed by Lee-Kwang et al. [15] was chosen to be incorporated into the framework.

$$S(\tilde{A}, \widetilde{B}) = \max_{x \in X} \min(\mu_{\tilde{A}}(x), \mu_{\widetilde{B}}(x)). \tag{2.47}$$

This measure is a special case of Enta's similarity that uses the minimum (Gödel's) $T$-norm (2.20) to implement the intersection operation [27].

*Note 1* There is also a concept of similarity as a type of fuzzy relation. It is an application of fuzzy set theory to describe similarity between any two elements from a given set. As such, it significantly differs from the idea of measuring the similarity of two fuzzy sets described in this section.

---

## 2.3 Linguistic Variables

When modeling qualitative data or discretizing quantitative data, some standardization of their description is usually needed. It means unified labels for their values and often also established ordering. This type of standardization can be handled by linguistic variables. A linguistic variable is a variable whose values are fuzzy sets that represent concepts from natural languages. More formally, a linguistic variable is defined as a quintuple [17] consisting of:

$V$ The name of the variable

$X$ The universe of discourse

$T(X)$ A set of terms that serve as labels for different values of the variable

$G$ A grammar that generates the items of $T(X)$

$M$ A semantic rule assigning meaning to items in $T(X)$ by mapping it to a fuzzy set membership function

The fuzzy sets accompanied by the labels define a fuzzy linguistic scale.

**Example** In the context of this book, we might consider a linguistic variable called *sentiment*. This variable is defined in a universe of polarity

values from a $[-1, 1]$ interval. The variable can have one of the three basic values labeled as *positive*, *neutral*, and *negative*. A corresponding fuzzy set for each of these three labels is defined. Moreover, it is possible to introduce grammar rules that would allow us to create new values. For example, one can add the adverb *very* to any label that corresponds to the operation $\mu_{(\text{very } \tilde{A})}(x) = \mu_{\tilde{A}}(x)^2$. A suitable linguistic scale can then be as follows (the triplets describe the corresponding TFNs), see Fig. 2.2:

- Extremely negative $(-1, -1, -0.75)$
- Rather negative $(-0.9, -0.75, -0.25)$
- Negative to neutral $(-0.5, -0.25, 0)$
- Neutral $(-0.25, 0, 0.25)$
- Neutral to positive $(0, 0.25, 0.5)$
- Rather positive $(0.25, 0.75, 0.9)$
- Extremely positive $(0.75, 1, 1)$

There are no strict rules for establishing the linguistic scale. In general, it is necessary to decide for a number of labels which enables to distinguish between two values which are qualitatively considerably different and which is, on the other hand, as low as possible. In addition, the labels should be clearly understood by the domain expert providing the evaluation. In order to create a high-quality linguistic variable, the quality indicators introduced in [1] should be met. These criteria are compatible with the Ruspini partition [21] (Fig. 2.6).

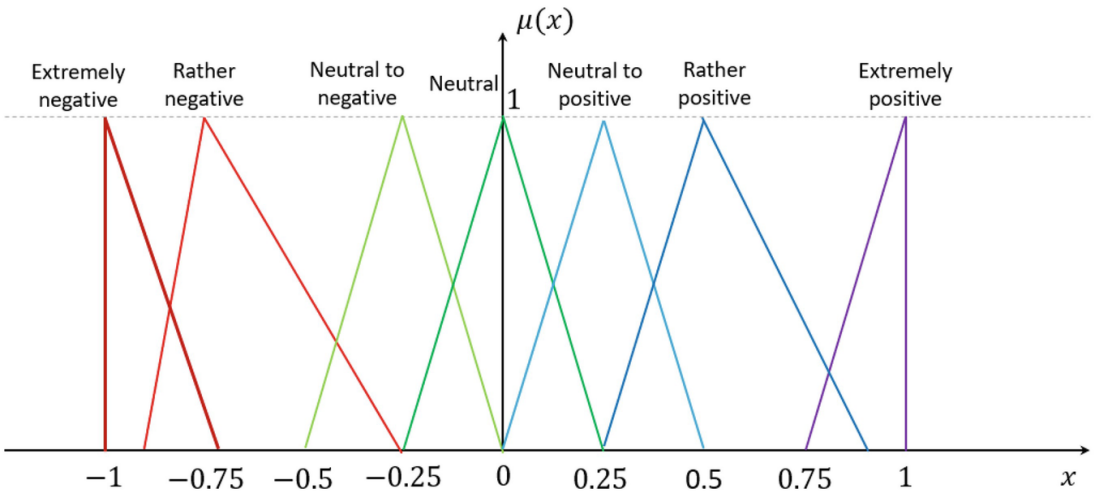


***Fig. 2.6*** Example of the linguistic scale with seven grades to assess the satisfaction

## 2.4 Estimating Membership Functions

One of the most challenging tasks when working with fuzzy sets is formalizing a membership function which properly describes a given

concept or a set of membership functions describing possible values of a linguistic variable. This topic is analyzed in great detail in [17]. What follows is a high-level overview of the most important ideas. The problem of membership function estimation can be approached from two perspectives: data-driven and user-driven (sometimes also called expert-driven).

### 2.4.1 Data-Driven Approaches

Data-driven approaches require a large dataset of sample objects from the domain of interest. If such a dataset is available, some type of automated procedure can be applied to estimate membership functions. Generally, equi-length, equi-depth, and equi-log are approaches to discretizing domain of a considered attribute into classical partitions [2, 7]. As for fuzzy partitions, it is possible to adopt these approaches supported by statistical measures such as data distribution and arithmetic mean by allowing overlapping between neighboring labels [22].

Fuzzy clustering is another approach. Similarly to the difference between fuzzy sets and crisp sets, fuzzy clustering methods perform what is sometimes referred to as *soft* clustering. Instead of strictly assigning each data point to a single cluster, fuzzy clustering allows a data point to be assigned to multiple clusters with different membership degrees for situations where the task is to determine the membership functions to neighboring fuzzy sets in a linguistic variable.

### 2.4.2 User-Driven Approaches

User-driven approaches rely on human judgment when estimating membership functions. This usually means asking one or more experts to determine the membership degree of several sample objects. This group of approaches can be divided into two subgroups: horizontal and vertical approaches.

Horizontal procedures determine the membership functions of individual data points. An example of such procedure is asking multiple experts the binary question: “Does value $x$ belong to the concept $A$?” The membership degree of $x$ in $A$ is calculated as the ratio of the number of positive answers to the number of all the answers.

Vertical procedures are based on the aforementioned concept of $\alpha$-cuts. In an implementation of a vertical procedure, one can ask an expert to select all elements that belong to some concept $A$ with a membership degree of $\alpha$ or higher. By asking this question for multiple values of $\alpha$, it is possible to approximate the membership function.

It is important to emphasize that vertical user-driven approaches, horizontal user-driven approaches, and data-driven approaches are not mutually exclusive, quite the contrary. Using multiple estimation

procedures has the added benefit of allowing the analyst to evaluate consistency.

Even a detailed and cautious construction of a linguistic variable is not immune to inconsistency. Hence, it is advisable to check whether the constructed linguistic variable meets the quality criteria of individual fuzzy sets and the linguistic variable proposed in [1].

# References


1. Alonso JM, Castiello C, Magdalena L, Mencar C (2021) Explainable fuzzy systems: paving the way from interpretable fuzzy systems to explainable AI systems. Springer, Cham [Crossref]
2. Aggarwal CC (2015) Data mining: the textbook. Springer, Cham [Crossref]
3. Beliakov G, Pradera A, Calvo T (2007) Aggregation functions: a guide for practitioners. Springer, Berlin Heidelberg
4. Beg I, Ashraf S (2009) Similarity measures for fuzzy sets. In: Applied and computational mathematics, vol 8(2)
5. Bector, CR, Chandra S (2005) Fuzzy mathematical programming and fuzzy matrix games. Springer, Berlin, vol 169
6. Bojadziev G, Bojadziev M (2007). Fuzzy logic for business, finance and management, 2nd edn. World Scientific, London [Crossref]
7. Bramer M (2020) Principles of data mining. Springer, London [Crossref]
8. Galindo J (2005) Fuzzy databases: modeling, design and implementation. IGI Global
9. Gottwald S (2015) Many-valued and fuzzy logics. In: Kacprzyk J, Pedrycz W (eds) Handbook of computational intelligence. Springer Berlin, Heidelberg, pp 7–29
10. Grabisch M, Marichal J L, Mesiar R, Pap E (2009) Aggregation functions. Cambridge University Press, Cambridge [Crossref]
11. Grabisch M (2003) Modelling data by the Choquet integral. In: Information fusion in data mining. Springer, Berlin, Heidelberg, pp 135–148 [Crossref]
12. Jones A, Kaufmann A, Zimmermann HJ (2012) Fuzzy sets theory and applications, vol 177. Springer, Berlin
13. Klement EP, Mesiar R, Pap E (2000) Triangular norms. Kluwer Academic Publishers, Dordrecht [Crossref]
14. Kóczy L, Hirota K (1993) Ordering, distance and closeness of fuzzy sets. Fuzzy Sets Syst 59(3):281–293 [Crossref]
15. Lee-Kwang H, Song YS, Lee KM (1994) Similarity measure between fuzzy sets and between elements. Fuzzy Sets Syst. 62(3):291–293 [Crossref]

16. Mishra AR, Jain D, Hooda D (2016) On fuzzy distance and induced fuzzy information measures. J Inform Optim Sci 37(2):193–211
17. Pedrycz W (2021) An introduction to computing with fuzzy sets: analysis, design, and applications. Springier, Cham
18. Inuiguchi, M, Ramık, J (2000) Possibilistic linear programming: a brief review of fuzzy mathematical programming and a comparison with stochastic programming in portfolio selection problem. Fuzzy Sets Syst 111(1):3–28
[Crossref]
19. Fiedler M, Nedoma J, Ramık J, Rohn J, Zimmermann K (2006) Linear optimization problems with inexact data. Springer, Berlin
20. Ramık J, Vlach M (2012) Generalized concavity in fuzzy optimization and decision analysis, vol 41. Springer, Berlin
21. Ruspini EH (1969) A new approach to clustering. Inform Control 15:22–32
[Crossref]
22. Tudorie C (2008) Qualifying objects in classical relational database querying. In: Galindo J (ed) Handbook of research on fuzzy information processing in databases, information science reference. Hershey, pp 218–245
23. van Leekwijck W, Kerre EE (1999) Defuzzification criteria and classification. Fuzzy Sets Syst 108:159–178
[Crossref]
24. Xuecheng L (1992) Entropy, distance measure and similarity measure of fuzzy sets and their relations. Fuzzy Sets Syst 52(3):305–318
[Crossref]
25. Zadeh LA (1965) Fuzzy sets. Inform Control 8(3):338–353. https://doi.org/1016/S0019-9958(65)90241-X
[Crossref]
26. Zadeh LA (1977) Fuzzy sets as a basis for a theory of possibility. Fuzzy Sets Syst 100:9–34. https://doi.org/1016/S0165-0114(99)80004-9
[Crossref]
27. Zwick R, Carlstein E, Budescu DV (1987) Measures of similarity among fuzzy concepts: a comparative analysis. Int J Approx Reason 1(2):221–242
[Crossref]

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2025
M. Švaňa et al., *From Text to Understanding*, Fuzzy Management Methods
https://doi.org/10.1007/978-3-032-00129-0_3

# 3. Machine Learning and Natural Language Processing

Miloš Švaňa[1] ✉, František Zapletal[1] and Miroslav Hudec[1]
(1) VSB – Technical University of Ostrava, Ostrava-Poruba, Czech Republic

## 3.1 Classification of Machine Learning Approaches

An agent is learning when it improves performance after exploring the world. When the agent is a computer, it is called Machine Learning (ML) [56].

Machine learning can be broadly defined as a set of methods for solving problems by learning from data rather than explicitly creating an algorithm to solve a given problem. ML methods are beneficial for problems where creating an algorithm would be labor-intensive and therefore costly, or in complex tasks it might be practically impossible [6]. This includes problems such as computer vision, natural language processing, time series prediction, and more recently protein folding [37]. ML itself lies at the intersection of many fields: statistics, probability theory, logic, optimization, operations research, and parallel or distributed computing.

Problems that can be solved with ML methods fall into several categories, each with unique properties [42, 58]:

supervised learning, unsupervised learning, and reinforcement learning. The next and relevant aspect is explaining the solution obtained [29].

### 3.1.1 Supervised Learning

The task of supervised learning is as follows. Given a training set of $N$ example input-output pairs, created by an unknown function $g$, the goal is to find a function $f$ capable of predicting a correct output, that is, approximate the true function $g$ [56]. Consequently, the function $f$ can then be used to predict unknown output for input not included in its training set. Two typical problems for this category are classification and regression.

Classification problems involve predicting a category label for a given input sample, where the output is discrete. In the most trivial case, the category is binary [9]. For example, a system can be developed to categorize an e-mail as spam or as a legitimate message. More complex classification problems involve selecting from multiple classes, e.g., classifying images into animals, scenery, portraits, or even problems when a text can belong to multiple categories at once (i.e., document belongs to statistic and economy).

The decision made by a classification model can be either hard or soft. A hard prediction means predicting a final label, while a soft prediction usually outputs probabilities of a given sample belonging to different categories. A soft classifier can be easily transformed into a hard classifier, for example, by selecting the label with the highest probability [2].

In contrast, regression means predicting a continuous value, such as the price of a financial asset. In this case, it is predicting the values of different assets at a given time point or the values of an asset over time [3]. Next, regression reveals the dependency and its direction, that is,

whether an increase in the value of an input attribute causes a decrease or increase in the output attribute [73].

Classification and regression share many similarities. The same algorithm can often be used to train both a classification model and a regression model (although some parameter changes are required). For example, classification and regression can be integrated into one model using ordinal sums of conjunctive and disjunctive functions, handling classification into binary classes (e.g., *Yes* and *No*) and intermediate class *Maybe* with keeping the intensity of inclination or using it as a threshold between classes [30].

An important limitation is that many ML methods assume similar probability distributions for training data and data on which the model is used. The model may fail to make a good prediction if the distributions are different, a phenomenon called *concept drift* [35].

The models are evaluated by the quality of predicting the output $f$ compared to the expected ground truth $g$. Different metrics are used for classification and regression tasks. The metrics for classification are *accuracy, precision, recall, F1*. More metrics are needed because classes can be imbalanced, or false negative and false positive might exhibit different cost [71]. An example of a regression metric is *mean square error*, i.e., the average of the squares of differences between the predicted value and the expected value [54].

Besides, more advanced procedures have been designed to uncover deeper issues. For example, *k-fold* cross validation [58] is useful to identify overfitting. And sometimes a third *validation set* can be introduced in addition to the training and test sets to properly validate the model [75].

Despite quality metrics, the cost to achieve labeled data should also be considered. It might be tedious work for

domain experts to evaluate input data and assign expected outputs [42, 56]. In this case, the option is unsupervised learning or merging supervised learning with unsupervised learning.

### 3.1.2 Unsupervised Learning

Unsupervised learning methods are used when the goal is not to find a mapping from a given input to an output, but to discover hidden patterns in an explored dataset. Often, there is no output variable included in the available data (or it is costly to obtain output behavior for all entities considered). Common unsupervised learning problems are clustering, dimensionality reduction, and associative rules [3].

Clustering can be defined as the process of segmenting a set of data points (entities such as customers, documents, or sentiments) into several groups or clusters. The problem is somewhat similar to classification. However, unlike class labels, the cluster centers and number of clusters are not known beforehand [9]. Clustering can therefore be used to discover hidden patterns and the most suitable number of clusters and relationships in a dataset.

Similarly like in classification, crisp and flexible clusters can be distinguished. In the latter, the entity might partially belong to several clusters with the sum of intensities of belonging equal to 1. If the final solution requires assigning to one cluster, then the winner is the cluster with the highest intensity of belonging [32].

Clustering is generally divided into two categories: representative and hierarchical [9]. In the former, a number of $k$ clusters are selected in advance. Successively, the final positions are learned to achieve the suitable solution. The well-known methods are k-means, k-median, and k-medoids. The main challenge compared to classification is the lack of quality metrics [3]. The grant

truth or output class, as well as the right number of clusters, is not known. The optimization is based on the minimization of the sum of the distances of all entities to their respective closest cluster. As the value $k$ increases, the distances decrease. When the decrease of function starts to be marginal, one gets the relevant number of clusters. Hence, the name of the often-used evaluation approach is *the elbow method*.

The main problem with the representative clustering is covering clusters that are far from being hyper-spherical, i.e., clusters exhibit different shapes. The solution is to divide space into a large number of small hyper cubes and merge neighboring ones when their densities are above the threshold values. For this problem, a known algorithm DBSCAN [3] and its HDBSCAN extension [13] are often applied.

In hierarchical clustering, it begins with the number of clusters equal to the number of entities and in successive steps assigns entities into the cluster having the shortest distance, while it ends with one root cluster. Returning, the key number of clusters on different levels can be recognized. The problem of nonspherical clusters is not so high, as there are several methods on how to assign an entity into a cluster such as single-ling, average-link, and complete-link [3]

Dimension reduction is used to reduce the dimensionality of the feature or attribute space of a given set of entities. This action can be performed as a preprocessing step when solving other machine learning problems to reduce training time or improve model quality. Dimension reduction also enables visualization of high-dimensional data. An example is to detect less relevant features using Hopkins statistics [28]. The next approach is the reduction of the dimensionality by principal component analysis [24].

Generally speaking, the number of dimensions can be reduced either by selecting a subset of original features or by transforming the original data into a completely new vector space consisting of principal components.

Naturally, the question why is dimensional reduction relevant applies, as someone might argue that each feature brings unique information, like when so-called big data are considered. It should be kept in mind that the features might be significantly correlated, the less relevant feature might exhibit the highest variability, etc. The $L_p$ metric is proven to be sensitive to a higher number of characteristics [3]. This holds for all values of $p \in \mathbb{N}$, including the most used $p = 1$ for the Manhattan distance and $p = 2$ for the Euclidean distance.

### 3.1.3 Reinforcement Learning

The work [66] defines reinforcement learning as learning that considers what action to take next in a given situation based on a numerical reward or punishment signal. The model selects an action to perform in a given situation and is then rewarded or punished based on how good the action was according to the selected criterion. Over time, the model learns to select actions that lead to the maximization of its total reward. In other words, the model acts as an agent trying to achieve a certain goal in some environment.

In supervised learning, the learning process is passively observing examples of input/output data from a dataset to find the solution, e.g., patterns [56]. In many tasks, it is less expensive or tedious to provide signals than to label a large dataset. Generally, two main approaches are:

- **Model-based reinforcement learning**, where the agent uses a transition model to help interpret the reward signal and to make the decision on how to act
- **Model-free reinforcement learning**, where the agent neither knows nor learns a transition model of the

environment, instead it learns a more direct representation of how to behave

Reinforcement learning is used extensively in robotics, game-playing systems, text summarization, and question answering [12]. Next, a social media analysis framework was developed to measure public sentiment on the energy transition [65], finding a broad public support for the transition to green energy.

### 3.1.4 Hybrid Approaches

There are problems whose specific characteristics combine the aspects of multiple problem classes or lack of a sufficient set of labeled data for learning. Two important categories of hybrid problems are *semi-supervised* and *self-supervised* learning.

Semi-supervised learning is conceptually situated between supervised and unsupervised learning. It handles a larger amount of unlabeled data in combination with a smaller set of labeled data for training [68]. Therefore, the problem lies at the intersection of classification and clustering. The inferred clusters from the unlabeled data significantly help to assign class boundaries.

A necessary condition is that the underlying marginal data distribution $p(x)$ in the input space contains information about the posterior distribution $p(y|x)$, where $x$ is the input vector and $y$ is the output. When this is the case, unlabeled data are used to gain information about $p(x)$ and consequently about the posterior distribution [79].

Self-supervised learning refers to the construction of an alternative supervised learning task that then finds a solution to some unsupervised learning task as a side effect [44]. Self-supervised learning belongs to unsupervised learning since there is no manual label involved. Unsupervised learning detects specific data patterns, such as clusters, while self-supervised learning aims at

recovering labels from a co-occurring input to related information.

In self-supervised NLP learning focused on word embedding, Continuous Bag of Words (CBOW) aims to predict input tokens based on context tokens, while Skip-Gram aims to predict context tokens based on input tokens [47]. Next, auto-regressive models are considered as *Bayes net structure*. The objective of auto-regressive language modeling is maximizing the likelihood [72]. In this way, it can perform complex tasks such as text summarization or question answering.

---

## 3.2 Machine Learning Methods

This section covers machine learning methods utilized as a part of the proposed framework, namely the Naive Bayes used for sentiment analysis, and HDBSCAN and neural networks, both used for candidate topic modeling.

### 3.2.1 Naive Bayes

Naive Bayes is a classification method explored, for example, in [78]. As the name suggests, this method is derived from the Bayes theorem, a well-known part of probability theory. During inference, the goal of the Naive Bayes classifier is to estimate the probability that an example $E$ belongs to class $C$ [9]:

$$p(C|E) = \frac{p(E|C)p(C)}{p(E)}. \tag{3.1}$$

where $p(C|E)$ is a conditional probability that example $E$ belongs to class $C$, while $p(E|C)$ is a conditional probability that for given $C$, $E$ belongs to that class, $p(E)$ and $p(C)$ are prior probabilities.

Considering binary classification, the simplest type of classification problems, the decision rule of a Naive Bayes classifier can be defined as

$$f_{nb}(E) = \frac{p(C=1|E)}{p(C=0|E)} >= 1. \quad (3.2)$$

If this inequality holds, example $E$ is labeled as 1. Otherwise, it is labeled as 0.

The term *naive* expresses an assumption that features that characterize each example are mutually independent, that is, there is no correlation between them. Although this is an unrealistic assumption, when a large amount of observation is available, the Naive Bayes classifier achieves competitive results in various classification tasks, including sentiment analysis. Under this assumption, $f_{nb}$ can be calculated as

$$f_{nb}(E) = \frac{p(C=1|E)}{p(C=0|E)} \prod_{i=1}^{M} \frac{p(x_i|C=1|E)}{p(x_i|C=0|E)}, \quad (3.3)$$

where $x_i$ represents the value of feature $i$, for example, $E$, with $M$ being the total number of features. All the probabilities in the equation can be learned from the training set. This probability estimation lies at the core of the training procedure.

One weak point is when not a single combination of input observations explains the output feature. Due to the multiplication of independent elementary conditional probabilities, the resulting probability gets the value 0 [3]. To avoid this problem, it is recommended to apply the Laplace smoothing parameter which adds a small probability to this occurrence.

### 3.2.2 HDBSCAN

Density-Based Spatial Clustering of Applications with Noise or DBSCAN for short uses the concept of density to assign cluster labels to different data points [59]. The density $\tau$ is defined as the number of data points within the radius $\epsilon$. The exact value of these parameters should be defined by the user, but it should reflect the data distribution to avoid that isolated values join clusters [3, 19].

Knowing the value of $\epsilon$, each observation in the dataset can be categorized as:

- A **core point** if at least $\tau$ data points fall within radius $\epsilon$ around the point
- A **border point** if the point is not a core point, but lies within the radius of $\epsilon$ from a core point or
- A **noise point** if the point is neither a core point nor a border point

After categorizing each point, the data can be clustered by performing two steps:

- Create a new cluster for each core point or a group of connected core points. The core points are connected if their mutual distance is less than $\epsilon$.
- Assign each border point to a cluster corresponding to the nearest core point.

The HDBSCAN method is an extension of DBSCAN capable of creating hierarchical tree-like structures of clusters and subclusters. Originally proposed by [13], HDBSCAN constructs a mutual reachability graph with vertices representing individual points in the dataset and edges with weights determined by the mutual reachability distance between the two points. This distance is calculated as

$$d_{mreach}(x_p, x_q) = \max\{d_{core}(x_p), d_{core}(x_q), d(x_p, x_q)\}, \quad (3.4)$$

where $d_{core}(x_p)$ is the distance of point $x_p$ to its core point and $d(x_p, x_q)$ is the mutual distance of the two points.

After constructing the graph, the noise points are eliminated by removing the edges with weights greater than $\epsilon$. Clusters are then created by iteratively increasing a weight threshold value and grouping points whose mutual reachability distance is less than this threshold.

### 3.2.3 Neural Networks and Transformers

The seminal paper [51] introduced a simplified computational model of a neuron that is now known as *perceptron*. Given a set of $M$ input values, the output of a perceptron cell can be calculated as

$$y = \phi(\sum_{i}^{M} x_i w_i + b), \tag{3.5}$$

where $x_i$ is the value of the $i$-th input variable, $w_i$ is its weight, $b$ is bias or noise, and $\phi$ is a so-called activation function.

A fundamental neural network consists of the following core elements [27]:

- A set of synapses or connecting links which is characterized by a weight of input feature
- An adder for summing the input weighted signals
- An activation function for neuron to proceed signal further
- An external bias (or noise) which increases or decreases input of the activation function

The exact shape of the activation function is problem-dependent. For example, a sigmoid function is often used in classification tasks due to its differentiability [38].

In order to perform a certain task, the weights of all neuron cells in the network must be “trained” on a training set that contains both the values of the input variable and the corresponding expected output. This training is usually implemented as a variant of the backpropagation algorithm [61]. It can be summarized in the following series of steps:

1. Set all weights to a random value.

2. Calculate the output of the neural network for all observations in the training set.

3. Calculate the error, i.e., the difference between the output provided by the neural network and the

expected ground truth.

4. Use a backpropagation mechanism to determine the "contribution" of each weight to the total error, and slightly adjust each weight to lower the error. Then go back to step 2.

The training procedure stops either after a predefined number of iterations or if the improvement in error compared to the previous iteration is below a certain threshold.

Over time, many different, more complex types of neural network have been proposed. For example, convolutional neural networks [41] perform well in tasks related to image processing, and recurrent neural networks [62] are being used to process sequential data.

One of the most recent innovations in the field of machine learning is transformers [10]. Their introduction changed the landscape of neural networks and enabled the development of state-of-the-art Large Language Models (known as LLMs) such as GPT. This neural network architecture was originally proposed by [69]. In this paper, transformers were applied to the problem of machine translation. This task was usually solved by some type of recurrent neural network. The main issue that the transformers were addressing is that recurrent neural networks are difficult to parallelize, significantly increasing training and inference times.

The transformer architecture is depicted in Fig. 3.1. It consists of two main parts: an encoder (left) and a decoder (right). In case of translation, the encoder takes the text in the original language as input. The input of the decoder part is the text in the target language that was produced so far. The network then predicts the next word to be added to the translation.

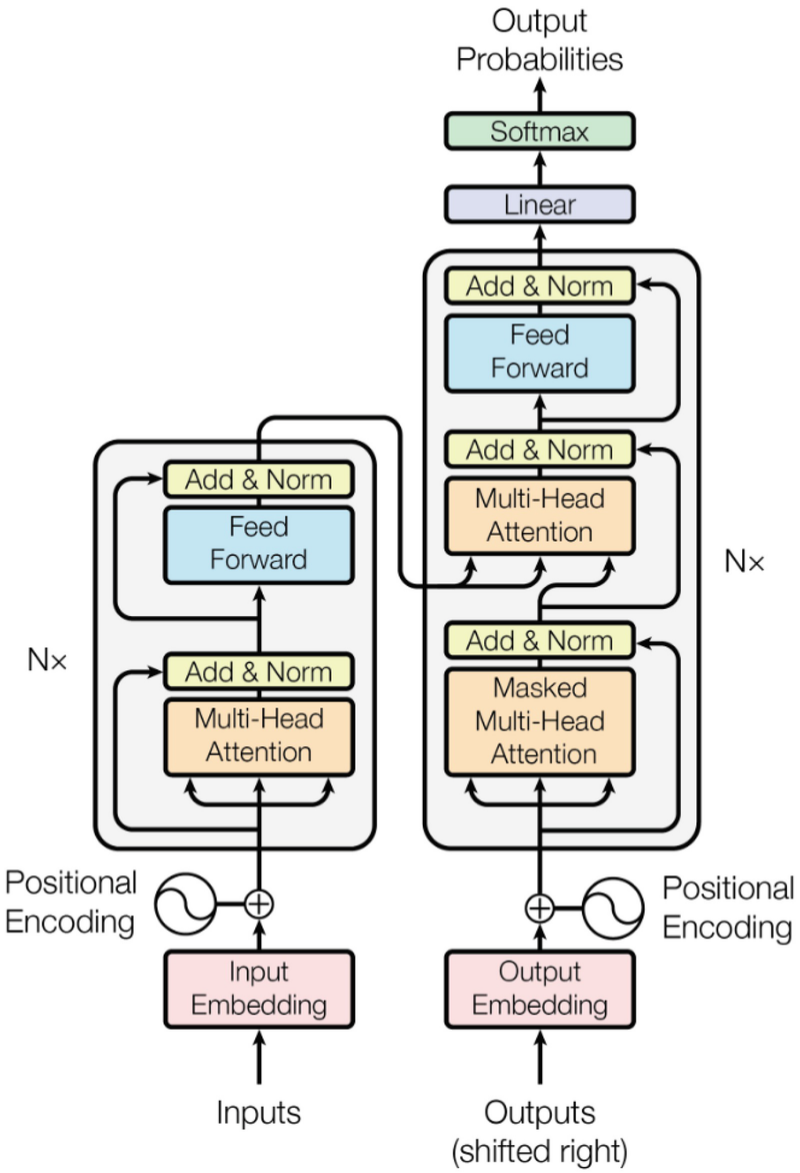


***Fig. 3.1*** A transformer neural network [69]

The input text is encoded using word embeddings. Unlike recurrent neural networks, a transformer does not implicitly carry information about the position of individual words in the text. The authors [69] solved this problem by adding a positional encoding to the word embeddings. This encoding is calculated as

$$PE(\text{pos}, 2i) = \sin\left(\text{pos} \,/\, 10000^{2i/d_{\text{pos}}}\right) \tag{3.6}$$

$$PE(\text{pos}, 2i+1) = \cos\left(\text{pos} \,/\, 10000^{2i/d_{\text{pos}}}\right). \tag{3.7}$$

While 3.6 is used for even dimensions in the input embedding, 3.7 is used for odd dimensions. The position of the word in the input text is defined by *pos*, *i* is the dimension of the embedding, and $d_{pos}$ is the total number of dimensions of the embedding.

The most important component of the transformer architecture is the *attention* mechanism. Attention produces a matrix that describes the mutual importance of input words given to the task. Transformers use multihead attention, so multiple different attention matrices are produced for a given input. Attention requires three inputs: $Q$ (query), $K$ (key), and $V$ (value). In the encoder part, all inputs are derived from the output of the previous layer (the transformer can have multiple layers, each consisting of an attention part and a feed-forward part) or from the original text embeddings. In the decoder part, the query is derived from the output of the previous layer or the translated text embeddings, but the key and the value are taken from the encoder.

Then the output of the attention layer is calculated as

$$\begin{aligned} \text{MultiHead}(Q,K,V) &= \text{Concat}(\text{head}_1,\text{head}_2,\ldots,\text{head}_n)W^O \\ \text{head}_i &= \text{Attention}(QW_i^Q, KW_i^K, KW_i^K) \\ \text{Attention}(Q,K,V) &= \text{softmax}\Big(\frac{QK^T}{\sqrt{d_k}}\Big)V, \end{aligned} \tag{3.8}$$

where $W$'s are weight matrices learned during training and $d_k$ is the number of input dimensions.

---

## 3.3 Natural Language Processing

Natural language processing or NLP for short lies at the intersection of artificial intelligence, computational intelligence (such as fuzzy logic), and linguistics. It deals with understanding and generating natural (human) language using computers, usually by machine learning

methods (e.g., the combination of NLP and ML is used for the classification of emails into spam/non-spam, language translators, information summarization, or question answering [40]). It is also related to the field of information retrieval, whose main goal is to use various methods of computer science to support text indexing and searching.

NLP covers a wide range of tasks. The work [53] distinguishes between low- and high-level NLP tasks and offers several examples for each category. Low-level NLP tasks include:

- Sentence boundary detection
- Tokenization, i.e., detection of individual tokens such as words
- Part-of-speech tagging of individual words
- Morphological decomposition such finding the basic dictionary form of a word (lemmatization)
- Chunking, i.e., detecting syntactically meaningful groups of words such as noun or verb phrases
- Other types of text segmentation

High-level NLP tasks are built on top of low-level tasks. They include:

- Spelling and grammar error detection and recovery
- Named entity recognition, i.e., identification of people, places, or specific types of objects
- Derivation of other word forms
- Detection of synonyms
- Word sense disambiguation, i.e., determining the exact meaning of homonyms
- Extraction of relationships between entities
- Machine translation
- Question answering
- Text summarization
- Topic modeling
- Sentiment analysis

The next concept is computing with words. It is also considered a low-level NLP, where it supports relaxing classical database queries, explaining flexible functional dependencies, and the like. An example is emphasizing regions where *most young people have rather negative opinions* or variants thereof of summarized sentences [31, 49] or to explain the correlation linguistically [50]. In the computing with words approach, fuzzy logic, fuzzy sets, and logic aggregation functions solve the task.

It is possible to distinguish between three phases of the development of NLP methods. In the first phase, most of the methods were symbolic and rule-based. Knowledge obtained from textual data was often represented in the form of a conceptual ontology. Following the work of [14], language processing tools were based on formal grammars, which are still used when dealing with programming languages and creating compilers or interpreters. Another important tool developed and still used today are regular expressions [64].

Phase 2 was data-oriented. It was dominated by statistical and classic machine learning methods such as the Naive Bayes described in Sect. 3.2.1. Other popular methods used in NLP during this phase included Support Vector Machines [15] or Hidden Markov Models [52].

We are currently in phase 3, which can also be called neural NLP or deep learning NLP. Deep neural networks became the go-to method for most NLP tasks [76]. One of the first breakthroughs that led to this paradigm shift was the introduction and popularization of word embedding models such as Word2Vec [48]. Word embeddings are dense vectors that represent individual words. They can capture word semantics in the sense that words with similar meanings are relatively close to each other in the embedding space. More recently, transformers discussed in Sect. 3.2.3 have become the dominant neural network

architecture in NLP. Their further development led to the creation of generative LLM and tools such as ChatGPT. A single large language model can perform a wide variety of NLP tasks, such as text summarization, question answering, content generation, or classification, without additional training [11]. When performing these tasks, generative models use the knowledge contained in the training corpora that was encoded into the model weights during training. For example, if a model is trained on a Wikipedia corpus, it can use the factual knowledge contained in this dataset to generate answers to factual questions. The work [26] suggests that LLM can encode even temporal and spatial information.

As already indicated, topic modeling and sentiment analysis are key components of the proposed social network analysis framework. Therefore, they are described in detail in the following subsections.

---

## 3.4 Topic Modeling

Topic modeling has become a vital tool for discovering hidden semantic structures in large amounts of texts. Through various ranges of methods, the main themes or topics are uncovered without the need to read the text line by line. This can help, for example, to summarize an extreme quantity of documents efficiently, for recommendation systems, or smart search. Although different authors phrase the definition differently or focus on different aspects, the basic ideas are consistent across the literature. For example, [39] offers two differently phrased definitions: Topic modeling is (1) a “statistical technique for revealing the underlying semantic structure in a large collection of documents” and (2) a “group of algorithms that reveal, discover, and annotate thematic structure in a collection of documents.”

The work [17] depicts the process of topic modeling graphically shown in Fig. 3.2. It has begun with a set of documents, e.g., books, and uses a topic model to identify topics that are contained in these documents. Topics are usually represented as sets of relevant words. After identifying the prevalent topics, it is possible to assign each document to one or more topics with different probabilities. In this sense, topic modeling can also be understood as a dimension reduction technique [39].

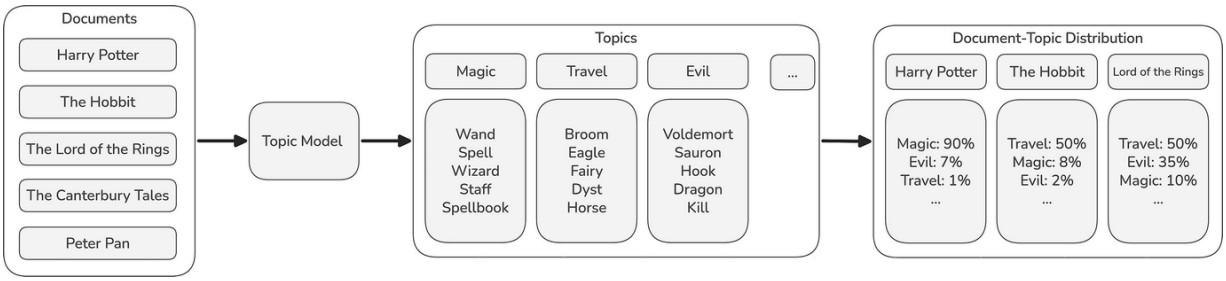


***Fig. 3.2*** Graphical depiction of the topic modeling process [17]

There is an abundance of topic modeling techniques. To understand the landscape, literature review articles provide different points of view. The work [17] explores the evolution of topic modeling over time, while [70] categorizes topic modeling methods based on their applicability to various types of input data and user intentions. Next, the study [39] sorts the topic modeling methods into two categories based on the underlying mathematical tools: (1) methods that rely on statistical properties of the input dataset and (2) algebraic methods usually built on top of matrix factorization techniques. Finally, the study [57] specifically focuses on methods for modeling short-text topics. This category is especially relevant, as the primary input for the proposed framework is social media posts. Compared to other types of content, such as books or scientific papers, these texts tend to be relatively short.

Two classic topic modeling approaches mentioned in the aforementioned reviews are Latent Semantic Analysis (LSA) and Latent Dirichlet Allocation (LDA). Introduced by [18], LSA was originally intended as a method to index and retrieve documents from large databases. It applies singular-value decomposition to a simple word-document matrix. In contrast, LDA [8] is a probabilistic model that models topics using the Dirichlet distribution, which is a multidimensional generalization of a beta distribution. In the context of LDA, each dimension of the Dirichlet distribution represents one topic. The work [39] compared both LSA and LDA in two standard textual datasets. In terms of topic coherence, LDA provided better results. The aforementioned literature review papers then list multiple LDA modifications suitable for different use cases, including correlation topic models, dynamic topic models for analyzing the evolution of topics over time, or hierarchical topic models capable of creating topic hierarchies.

The study [17] recognizes nonnegative matrix factorization (NMF) as another milestone in the development of topic modeling methods. The goal is to represent a word-document matrix as a product of two much smaller word-topic and document-topic matrices. NMF is an NP-hard problem, and as such various approximations have to be used.

Recently developed topic modeling approaches usually combine traditional methods with other NLP techniques, mainly word embedding models such as Word2Vec [48] or BERT [20, 25]. State-of-the-art generative transformer models, such as those of the GPT family [11], can also perform topic modeling.

Looking specifically at topic models intended for use with short text, the work [57] categorizes the methods into three groups:

- *Dirichlet multinomial mixture (DMM) methods* assuming that each document belongs to one single topic. This assumption can be problematic for long documents, such as books, but holds for many short-text documents.
- *Global word co-occurrence methods* assuming that co-occurring words are related to the same topic
- *Self-aggregation methods* address the issue of short text by aggregating documents into larger pseudo-documents. For example, the work [77] aggregates Twitter posts by their authors.

Next, the study [57] compares the quality of the extracted topics using different topic modeling methods on several datasets. The best-performing model for the Twitter dataset in terms of accuracy and coherence is LF-DMM [55].

### 3.4.1 Topic Model Evaluation

The above analysis of different topic modeling methods created several statements about the quality and performance of different models. Several ways to measure the quality of the topic model exist. The study [17] distinguishes three main groups:

1. *Coverage* measures. These measures require some source of ground truth, e.g., each document needs to be assigned to a topic by a human. An example of a metric belonging to this category is *topic recall* which can be calculated as the percentage of ground truth topics discovered by a specific model.

2. *Coherence* measures, such as pointwise mutual information, take into account the noise and intelligibility of the topics discovered.

3. *Qualitative human evaluation*. Neither coverage nor coherence measures can holistically evaluate whether a

topic model “makes sense” to a human. An approach to evaluating topics by a human is to ask the evaluator to compare similar topics generated by different methods.

Since there is no ground truth for the problem addressed in this study, the topic diversity and topic coherence measures were selected for further evaluation. Topic diversity is a simple ratio of unique words in each topic, i.e., words that do not occur in the description of other topics.

Topic coherence is calculated as normalized pointwise mutual information:

$$NPMI(w_i, w_j) = \frac{PMI(w_i,w_j)}{-\log_2 p(w_i,w_j)}, \tag{3.9}$$

where

$$PMI(w_i, w_j) = \log_2 \frac{p(w_i,w_j)}{p(w_i)p(w_j)} \tag{3.10}$$

The expression $p(w_i, w_j)$ represents the probability of the co-occurrence of words $w_i$ and $w_j$ in a single document, while $p(w_i)$ and $p(w_j)$ represent the occurrence probability of words $w_i$ and $w_j$ individually.

### 3.4.2 Choosing an Appropriate Model

A high number of different topic modeling approaches exist. Selecting a suitable method for a given use case can be a challenging task. Figure 3.3 [17] shows several factors that must be considered when selecting a topic modeling approach. To this end, [70] provides a decision tree for finding a method appropriate for a given situation. This tree is depicted in Fig. 3.4. Moreover, besides choosing an appropriate method, one also has to consider different parameters that have to be set by a method user. For example, many methods require the user to choose the number of topics to discover.

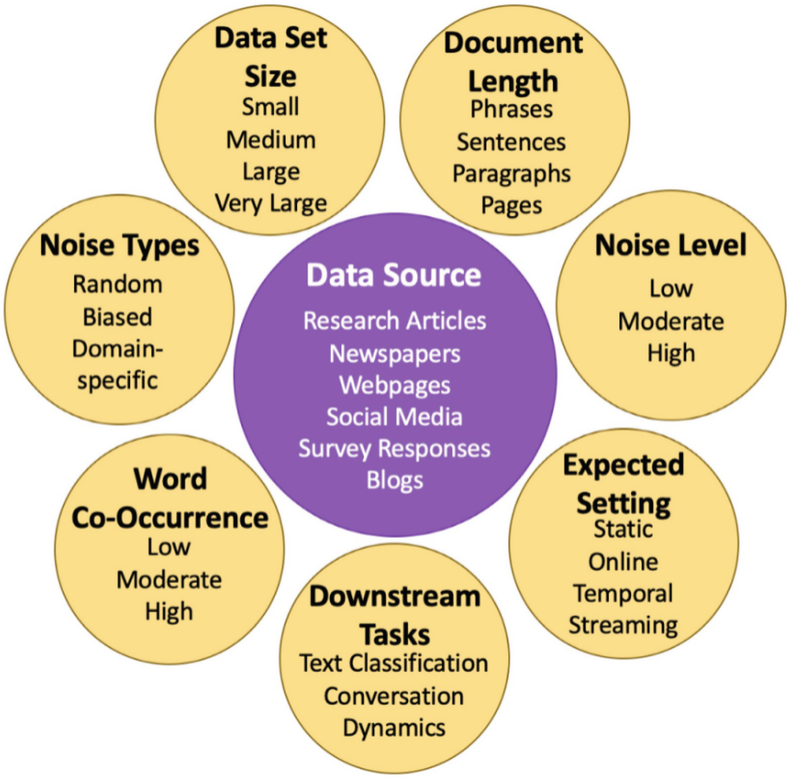


***Fig. 3.3*** Factors to consider when selecting a topic model [17]

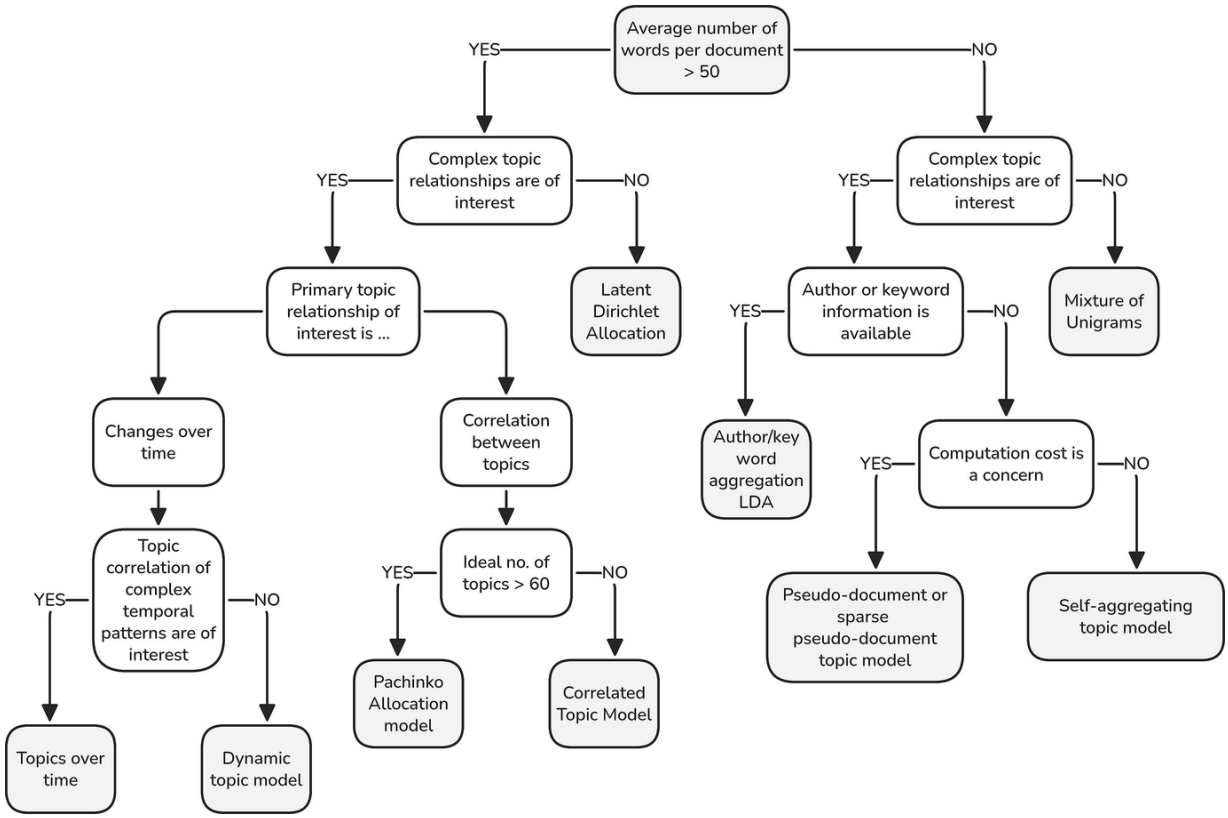
Average number of words per document > 50
YES
NO
Complex topic relationships are of interest
Latent Dirichlet Allocation
Primary topic relationship of interest is ...
Changes over time
Topic correlation of complex temporal patterns are of interest
Topics over time
Dynamic topic model
Correlation between topics
Ideal no. of topics > 60
Pachinko Allocation model
Correlated Topic Model
Complex topic relationships are of interest
Mixture of Unigrams
Author or keyword information is available
Author/key word aggregation LDA
Computation cost is a concern
Pseudo-document or sparse pseudo-document topic model
Self-aggregating topic model

***Fig. 3.4*** A decision tree for selecting a topic model [70]

A potential solution to the issue of choosing a proper method is provided by [7] who suggest aggregating multiple topic models with different configurations. Their experiments show a significant improvement in coherence.

Since a detailed examination of different topic modeling methods is not the main goal of this work, only two well-known methods were selected for further consideration: LDA and BERTopic. LDA represents a more traditional approach to topic modeling, while BERTopic employs the latest advances in neural network architectures. This contrast is useful for examining the benefits that modern deep learning methods bring to topic modeling.

### 3.4.3 Latent Dirichlet Allocation

The definition of LDA is as follows [8]:

> Latent Dirichlet allocation (LDA) is a generative probabilistic model of a corpus. The basic idea is that documents are represented as random mixtures over latent topics, where each topic is characterized by a distribution over words.

LDA assumes that each document in a corpus is generated by the following process:

- Choose a number of words in the document $N$ from a Poisson distribution.
- Choose a mixture of topics $\theta$ from a Dirichlet distribution. A mixture of topics is a list of real numbers from a $[0, 1]$ interval whose values sum to 1, for example, $(0.5, 0.3, 0.2)$. This specific list defines a given document as a mixture of topic 1, topic 2, and topic 3 in a 5:3:2 ratio.
- Repeat $N$ times:
  - Choose a topic using the topic mixture as probabilities.

– Choose a word $w_n$ while respecting the probability distribution $p(w_n|z_n, \beta)$ with $z_n$ being a topic selected in the previous step and $\beta$ being a matrix containing the probabilities of words belonging to specific topics.

During topic modeling, the task is inverted. Instead of generating a new document given a topic distribution and a word-topic probability matrix, the goal is to determine a topic mixture of a document or a set of documents given observed words. The study [8] claims that this task is intractable and mentions several approximation methods.

Before using LDA to perform topic inference, both the word-topic matrix $\beta$ and the $\alpha$ parameter of the Dirichlet distribution of topic mixtures must be estimated from available training data. Given a set of $M$ documents in the training set, this is done by finding such $\alpha$ and $\beta$ that maximize the likelihood of observing words of each document:

$$\sum_{d=1}^{M} log\ p(\mathbf{w}_d|\alpha, \beta). \tag{3.11}$$

An important practical consideration when using LDA is that the user has to choose the number of topics to identify, usually denoted $k$. This parameter cannot be derived from the training data.

### 3.4.4 BERTopic

Originally introduced by [25], BERTopic approaches the problem of topic modeling as a clustering task. The procedure consists of the following steps:

1. **Create an embedding for each document**. One of the downsides of traditional methods, such as LDA, is that they represent documents in a bag-of-words fashion. This representation ignores both the order of words in a document and their semantic relationships. In contrast, BERTopic relies on an embedding representation. As the name suggests, BERTopic uses

embeddings based on Sentence-BERT [60], a modified version of BERT [20].

2. **Reduce embedding dimensionality**. Embeddings can have hundreds and even thousands of dimensions. The study [25] mentions two problems caused by high dimensionality in clustering tasks. First, according to [4], the difference between the distance to the closest point to a cluster center and the distance to the furthest point from a cluster center decreases. Second, empirical research shows that a lower number of dimensions tend to lead to better performance in terms of time and quality. Therefore, BERTopic uses UMAP [46] to reduce the number of dimensions.

3. **Use a clustering algorithm to create document clusters**. BERTopic uses the HDBSCAN algorithm introduced in Sect. 3.2.2 with reduced document embeddings by UMAP as input. The user can set the number of clusters/topics to obtain from the algorithm. To provide the desired number of topics, small similar topics discovered by HDBSCAN are merged.

4. **Create topic representations**. In addition, BERTopic represents each topic as a list of words. An intuitive approach to selecting these words would be to use words closest to a cluster center. However, the work [25] proposes a modified TF-IDF (term frequency—inverse document frequency) score that should result in a more comprehensive representation. This metric is calculated as follows:

$$W_{t,c} = tf_{t,c}\, log(1 + \frac{A}{tf_t}), \tag{3.12}$$

where $tf_{t,c}$ represents the frequency of the term $t$ in the cluster $c$. During the calculation of the TF-IDF score, each cluster $c$ can be understood as a single document created by merging all documents belonging

to that cluster. $A$ is the average number of words per cluster, and $tf_t$ is the total frequency of the term $t$ in all clusters.

An important feature of BERTopic is that it can perform soft document labeling. Each document can be assigned to multiple topics with different probabilities. These topic probabilities can be used as weights when determining the overall sentiment toward each topic.

The issue of topic model selection in the context of the proposed social network data analysis framework is further addressed in Sect. 4.3.

---

## 3.5 Sentiment Analysis

Individuals and organizations make decisions that are to some extent based on the opinions of others. Surveys and focus groups are two examples of tools traditionally used by many organizations to gather and evaluate opinions of people belonging to some group of interest. These methods have undoubtedly their advantages. However, they can also be expensive and labor-intensive [23].

User-generated content published on the Internet, such as reviews, blog articles, or social media posts, represents an important alternative source of opinions. However, processing such large amounts of data manually is impractical. Tools automating this task had to be developed. Sentiment analysis (also called opinion mining) is one such tool.

The work [22] referring to [5] defines sentiment analysis as *an approach that uses Natural Language Processing (NLP) to extract, convert, and interpret opinions from a text and classify them into positive, negative, or natural sentiments*. The study [63] uses a much shorter definition: *sentiment analysis is a process of studying public opinion about an entity*.

Both definitions refer to the concept of *opinion*. The work [45] defines an opinion in the context of sentiment analysis as a quintuple consisting of:

- An object *o* towards which an opinion is being expressed
- An aspect *f* of this object that is being targeted
- Opinion orientation *oo*
- Author of the opinion *a*
- Time *t* at which the opinion was expressed

There are multiple ways of expressing the orientation of opinion. The simplest option is a categorical variable with two possible values: *positive* and *negative*. In some cases, a third *neutral* option is added. Alternatively, one can use various types of scale, such as a 5-star rating system [23].

Multiple sources, including [22, 23, 63, 74], distinguish between three levels of sentiment analysis:

- Document level, i.e., evaluating a given text document as a whole
- Sentence level
- Aspect level, i.e., focusing on individual aspects of an object, such as a laptop's performance and design

The same authors also broadly categorize the methods used to perform sentiment analysis into two main groups:

- **Lexicon-based**: Lexicon-based methods use predefined lexica of positive and negative words. One of the most popular sentiment lexicons is SentiWordNet.[1] Knowing the sentiment orientation of individual words, one can determine the overall sentiment of a sentence or a document by aggregating these individual word sentiments. One of the greatest advantages of lexicon-based methods is the relatively low labor requirements. However, there are also significant disadvantages. This group of methods is not suitable for dealing with issues such as slang or negation. There are also many words

that can be perceived both positively and negatively depending on context, e.g., “This vacuum sucks!”

- **Machine learning (ML)-based**: Alternatively, sentiment analysis can be approached as a classification or regression problem that can be solved by training a supervised machine learning model. Naive Bayes models and Support Vector Machines are among the most used ML methods in sentiment analysis. Generally speaking, ML methods provide better results. However, to train an ML model, some sort of ground truth has to be available. For some use cases, these labels are easily available. For example, many product reviews contain both text and a number of stars that the author gave to a product. However, quite often these labels have to be assigned manually by a human evaluator.

The work [22] examined the use of both methods in the literature. They found that of the 24 papers analyzed, seven used a lexicon-based method, ten used a machine learning method, and seven used a combination of both. Next, the study [21] claims that the literature published before their article suggests that lexicon-based methods perform the worst, machine learning methods the best, and their combination falls in the middle. However, their own experiments show that a combined approach outperforms both a lexicon-based method and a machine learning model.

The literature mentions several other types of sentiment analysis method. The work [74] discusses semi-supervised methods for sentiment analysis, an example of which is graph-based methods that use relationships between social network users and other entities to improve the results of sentiment analysis. Next, the work [67] focuses specifically on deep learning methods built on top of advanced neural network architectures with multiple layers. These include

convolutional neural networks, recurrent neural networks, and their various variants and combinations.

One of the main contributions of this work is the application of fuzzy modeling techniques to sentiment aggregation. Some authors have already applied models based on fuzzy logic or fuzzy set theory to sentiment analysis. For example, the authors [36] have developed a fuzzy expert system to determine the sentiment polarity of a document. The study [34] created a sentiment analysis method based on a fuzzy thesaurus where the relationships between words are fuzzy, with the strength of the relationship being determined by their relative distances within the documents where these words co-occur. However, in general, there seems to be little interest in fuzzy modeling among sentiment analysis researchers.

In terms of future challenges and opportunities, one of the most promising directions is multimodal sentiment analysis [74]. Social media users upload increasing amounts of non-textual content, including photos, videos, or audio recordings. Combining these types of data with textual information can significantly improve the capabilities of sentiment analysis systems.

Although not extensively, sentiment analysis has also been applied in the domain of smart cities. The study [43], for example, uses a Naive Bayes classifier to evaluate the sentiment of relevant tweets. They were able to use geographical information to visualize sentiment on a map. This technique can be especially useful in the domain of smart cities. The work [16] examined the sentiment contained in posts on social networks related to traffic in three Moroccan cities. Different local dialects presented a unique challenge.

The work [1] recognized many opportunities and challenges related to the use of sentiment analysis in smart city applications:

Future research is expected to address issues such as visual representation, multilingual audiovisual opinion mining, building usable peer-to-peer opinion mining tools for citizens, real-time sentiment analysis, multilingual reference corpora, recommendation algorithms, and automatic irony detection.

In comparison, the main contribution of this work in terms of sentiment analysis is the application of fuzzy modeling to capture uncertainty and the creation of a high-level framework capable of capturing sentiment related to various topics rather than focusing on specific issues.

Hence, the framework proposed in this work is based on existing sentiment analysis methods. As in the case of topic modeling methods, a detailed analysis of sentiment analysis methods is not the main goal of the work. Therefore, only three sentiment analysis methods are considered. They were selected mainly for their ease of use. Two of them are part of the TextBlob[2] library for Python, and one is available as a standalone VADER[3] library.

TextBlob is a Python library whose goal is to simplify various text analysis tasks, including sentiment analysis. The library offers two sentiment analysis methods: a lexicon-/pattern-based one and a pretrained Naive Bayes model. The former uses a polarity lexicon and structural patterns in the text to determine both the degree of sentiment polarity and the degree of subjectivity. Among other things, this method can consider negations or adverbs when determining the overall sentiment.

The Naive Bayes model is usually trained to classify reviews as positive or negative. Instead of providing degrees of polarity and intensity, this method returns the probability of a given text being positive and negative.

VADER (Valence Aware Dictionary and sEntiment Reasoner), proposed by [33], is a lexicon-/pattern-based sentiment analysis system designed specifically for social

networks. It considers elements such as negations or words written in capital letters. The output contains information about degrees of positivity, neutrality, and negativity, which is then aggregated into a single number that describes the overall sentiment.

# References


1. Ahmed KB, Radenski A, Bouhorma M, Ahmed MB (2016) Sentiment analysis for smart cities: state of the art and opportunities. In: Proceedings on the international conference on internet computing (ICOMP), pp 55–61
2. Alonso JM, Castiello C, Magdalena L, Mencar C (2021) Explainable fuzzy systems: paving the way from interpretable fuzzy systems to explainable AI systems. Springer, Cham
3. Aggarwal CC (2015) Data mining: the textbook. Springer, Cham
4. Aggarwal CC, Hinneburg A, Keim DA (2001) On the surprising behavior of distance metrics in high dimensional space. In: Van den Bussche J, Vianu V (eds) Database theory—ICDT 2001. Springer, Heidelberg, pp 420–434
5. Agarwal B, Mittal N, Bansal P, Garg S (2015) Sentiment analysis using common-sense and context information. Comput. Intell. Neurosci. 2015:715–730. https://doi.org/1155/2015/715730
6. Alpaydin E (2020) Introduction to machine learning. The MIT Press, Cambridge. https://api.semanticscholar.org/CorpusID:37364431
7. Blair SJ, Bi Y, Mulvenna MD (2016) Increasing topic coherence by aggregating topic models. In: Lehner F, Fteimi N (eds) Knowledge science, engineering and management. Springer, Cham, pp 69–81
8. Blei DM, Ng A, Jordan MI (2001) Latent Dirichlet allocation. J. Mach. Learn. Res. 3:993–1022. https://doi.org/5555/944919.944937
9. Bramer M (2020) Principles of data mining. Springer, London
10. Braşoveanu AM, Andonie R (2020) Visualizing transformers for NLP: a brief survey. In: 2020 24th IEEE international conference information visualisation (IV), pp 270–279
11. Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A, Agarwal S, Herbert-Voss A, Krueger G, Henighan T, Child R, Ramesh A, Ziegler D, Wu J, Winter C, Hesse C, Chen M, Sigler E, Litwin M, Gray S, Chess B, Clark J, Berner C, McCandlish S, Radford A, Sutskever I, Amodei D (2020) Language Models

are Few-Shot Learners. In: Larochelle H, Ranzato M, Hadsell R, Balcan M F, Lin H (eds) Advances in neural information processing systems. Curran Associates, pp 1877–1901

12. Botteghi N, Poel M, Brune C (2015) Unsupervised representation learning in deep reinforcement learning: a review. IEEE Control Syst. 45(2):26–68. https://doi.org/13140/RG.2.2.25102.43847

13. Campello RJGB, Moulavi D, Sander J (2013) Density-based clustering based on hierarchical density estimates. In: Pei J, Tseng V S, Cao L, Motoda H, Xu G (eds) Advances in knowledge discovery and data mining. Springer, Berlin, Heidelberg, pp 160–172

14. Chomsky N (1956) Three models for the description of language. IRE Trans Inform Theory 2(3):113–124

15. Cristianini N, Shawe-Taylor J (2000) An introduction to support vector machines and other kernel-based learning methods. Cambridge University Press, Cambridge. https://doi.org/1017/CBO9780511801389

16. Dahbi M, Saadane R, Mbarki S (2019) Social media sentiment monitoring in smart cities: an application to Moroccan dialects. In: Proceedings of the 4th international conference on smart city applications, Casablanca, pp 1–6. ACM, New York. https://doi.org/1145/3368756.3368997

17. Churchill R, Singh L (2022) The evolution of topic modeling. ACM Comput. Surv. 54(10s):1–35. https://doi.org/1145/3507900

18. Deerwester SC, Dumais ST, Landauer TK, Furnas GW, Harshman RA (1990) Indexing by latent semantic analysis. J Am Soc Inf Sci 41:391–407

19. Deng D (2020) DBSCAN clustering algorithm based on density. In: 2020 7th IEEE international forum on electrical engineering and automation (IFEEA). IEEE, pp 949–953

20. Devlin J, Chang MW, Lee K, Toutanova K (2019) BERT: pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805

21. Dhaoui C, Webster CM, Tan LP (2017) Social media sentiment analysis: lexicon versus machine learning. J Consum Market 34(6):480–488. https://doi.org/1108/JCM-03-2017-2141

22. Drus Z, Khalid H (2019) Sentiment analysis in social media and its application: systematic literature review. Proc Comput Sci 161:707–714. https://doi.org/1016/j.procs.2019.11.174

23. Farhadloo M, Rolland E (2016) Fundamentals of sentiment analysis and its applications. In: Pedrycz W, Chen S M (eds) Sentiment analysis and ontology engineering. Springer, Cham, pp 1–24

24. Gewers FL, Ferreira GR, Arruda HFD, Silva FN, Comin CH, Amancio DR, Costa LDF (2021) Principal component analysis: a natural approach to data exploration. ACM Comput Surv 54(4):1–34
25. Grootendorst M (2022) BERTopic: neural topic modeling with a class-based TF-IDF procedure. arXiv preprint arXiv:2203.05794
26. Gurnee W, Tegmark M (2023) Language models represent space and time. arXiv preprint arXiv:2310.02207
27. Haykin S (2009) Neural networks and learning machines, 3rd edn. Prentice Hall, New Jersey
28. He H (2025) Clustering algorithm based on Hopkins statistics and K-means. Multimedia Tools and Applications in press
29. Holzinger A, Kieseberg P, Weippl E, Tjoa AM (2018) Current advances, trends and challenges of machine learning and knowledge extraction: from machine learning to explainable AI. In: Springer lecture notes in computer science LNCS 11015. Springer, Cham
30. Hudec M, Mináriková E, Mesiar R, Saranti A, Holzinger A (2021) Classification by ordinal sums of conjunctive and disjunctive functions for explainable AI and interpretable machine learning solutions. Knowl.-Based Syst. 220:106916
31. Hudec M, Malovcová K, Trumic R, Rakovská E (2022) A new quality measure and visualization of the short-quantified sentences of natural language on maps—a case on COVID-19 data. Informatica, 33(2):321–342
32. Huang H, Shi W, Yang F, Feng Y, Zhang L, Liang X, Shi J, Cheng G, Huang J, Liu Z (2025) Fuzzy K-means clustering with reconstructed information. Int J Mach Learn Cybernet 16:43–53
33. Hutto C, Gilbert E (2014) VADER: a parsimonious rule-based model for sentiment analysis of social media text. In: Proceedings of the international AAAI conference on web and social media, vol 8(1), pp 216–225. https://doi.org/1609/icwsm.v8i1.14550
34. Ismail HM, Belkhouche B, Zaki N (2018) Semantic Twitter sentiment analysis based on a fuzzy thesaurus. Soft Comput 22(18):6011–6024. https://doi.org/1007/s00500-017-2994-8
35. Iwashita AS, Papa JP (2019) An overview on concept drift learning. IEEE Access 7:1532–1547. https://doi.org/1109/ACCESS.2018.2886026
36. Jefferson C, Liu H, Cocea M (2017) Fuzzy approach for sentiment analysis. In: 2017 IEEE international conference on fuzzy systems (FUZZ-IEEE),

Naples, Italy. IEEE, pp 1–6. https://doi.org/1109/FUZZ-IEEE.2017.8015577

37. Jumper J, Evans R, Pritzel A, Green T, Figurnov M, Ronneberger O, Tunyasuvunakool K, Bates R, Žídek A, Potapenko A, Bridgland A, Meyer C, Kohl SAA, Ballard AJ, Cowie A, Romera-Paredes B, Nikolov S, Jain R, Adler J, Back T, Petersen S, Reiman D, Clancy E, Zielinski M, Steinegger M, Pacholska M, Berghammer T, Bodenstein S, Silver D, Vinyals O, Senior AW, Kavukcuoglu K, Kohli P, Hassabis D (2021) Highly accurate protein structure prediction with AlphaFold. Nature 596(7873):583–589. https://doi.org/1038/s41586-021-03819-2

38. Keller JM, Liu D, Fogel DB (2016) Fundamentals of computational intelligence: neural networks, fuzzy systems, and evolutionary computation. John Wiley & Sons, Hoboken

39. Kherwa P, Bansal P (2018) Topic modeling: a comprehensive review. ICST Trans Scalable Inform Syst. https://doi.org/4108/eai.13-7-2018.159623

40. Khurana D, Koli A, Khatter K, Singh S (2023) Natural language processing: state of the art, current trends and challenges. Multimedia Tools Appl 82(3):3713–3744

41. Krizhevsky A, Sutskever I, Hinton GE (2017) ImageNet classification with deep convolutional neural networks. Commun ACM 60(6):84–90. https://doi.org/1145/3065386

42. Kwok JT, Zhou ZH, Xu L (2015) Machine learning. In: Kacprzyk J, Pedrycz W (eds) Springer handbook of computational intelligence. Springer handbooks. Springer, Berlin, Heidelberg

43. Li M, Ch'ng E, Chong A, See S (2016) The new eye of smart city: novel citizen sentiment analysis in Twitter. In: 2016 International conference on audio, language and image processing (ICALIP), Shanghai, China. IEEE, pp 557–562. https://doi.org/1109/ICALIP.2016.7846617

44. Liu X, Zhang F, Hou Z, Mian L, Wang Z, Zhang J, Tang, J (2021) Self-supervised learning: generative or contrastive. IEEE Trans Knowl Data Eng 35(1):857–876

45. Liu B (2010) Sentiment analysis and subjectivity. In: Handbook of natural language processing

46. McInnes L, Healy J, Melville J (2020) UMAP: uniform manifold approximation and projection for dimension reduction. arXiv preprint arXiv:1802.03426

47. Mikolov T, Sutskever I, Chen K, Corrado GS, Dean J (2013) Distributed representations of words and phrases and their compositionality. In: 33rd

international conference on neural information processing systems, pp 3111–3119

48. Mikolov T, Chen K, Corrado G, Dean J (2013) Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781

49. Mináriková E, Pisoni G, Molnár B, Skaftadottir H (2024) Business intelligence reporting by linguistic summaries for smart cities: a case on explaining bicycle sharing patterns. In 26th international conference on enterprise information systems, vol 2, pp 762–768

50. Mináriková E, Vučetić M, Hudec M (2023) The synergy of linear regression, fuzzy functional dependencies and linguistic summaries: a case of heat islands in Bratislava. In: 2023 IEEE XXIX international conference on information, communication and automation technologies (ICAT), pp 1–6

51. McCulloch W, Pitts W (1943) A logical calculus of ideas immanent in nervous activity. Bull Math Biophys 5:127–147

52. Mor B, Garhwal S, Loura A (2020) A systematic review of hidden Markov models and their applications. archives of computational methods in engineering. https://doi.org/1007/s11831-020-09422-4

53. Nadkarni PM, Ohno-Machado L, Chapman WW (2011) Natural language processing: an introduction. J Am Med Inform Assoc 18(5):544–551. https://doi.org/1136/amiajnl-2011-000464

54. Naidu G, Zuva T, Sibanda EM (2023) A review of evaluation metrics in machine learning algorithms. In: Silhavy R, Silhavy P (eds) Artificial intelligence application in networks and systems. CSOC 2023. Lecture notes in networks and systems, vol 724. Springer, Cham, pp 15–25

55. Nguyen DQ, Billingsley R, Du L, Johnson M (2018) Improving topic models with latent feature word representations. arXiv preprint arXiv:1810.06306

56. Norvig P, Russell S (2021) Artificial intelligence. a modern approach, 4th edn. Pearson Education, Harlow

57. Qiang J, Qian Z, Li Y, Yuan Y, Wu X (2022) Short text topic modeling techniques, applications, and performance: a survey. IEEE Trans Knowl Data Eng 34(3):1427–1445. https://doi.org/1109/TKDE.2020.2992485

58. Raschka S (2015) Python machine learning. Packt Publishing

59. Raschka S, Mirjalili V (2017) Python machine learning. Packt Publishing. https://books.google.fi/books?id=_plGDwAAQBAJ

60. Reimers N, Gurevych I (2019) Sentence-BERT: sentence embeddings using siamese BERT-networks. arXiv preprint arXiv:1908.10084

61. Rumelhart DE, Hinton GE, Williams RJ (1986) Learning representations by back-propagating errors. Nature 323:533–536. https://api.semanticscholar.org/CorpusID:205001834
62. Schmidt RM (2019) Recurrent neural networks (RNNs): a gentle introduction and overview. arXiv preprint arXiv:1912.05911
63. Sharma D, Sabharwal M, Goyal V, Vij M (2020) Sentiment analysis techniques for social media data: a review. In: Luhach AK, Kosa JA, Poonia RC, Gao XZ, Singh D (eds) First international conference on sustainable technologies for computational intelligence. Springer, Singapore, pp 75–90
64. Shannon CE, McCarthy J (1956) Automata studies. Princeton University Press, Princeton
65. Sun S, Guo Z, Li L, Zheng Z, Wang J, Anno S, Qian X (2025) Decoding public sentiments on energy transition in Japan's three major metropolitan areas: social media analysis using machine learning. J Cleaner Prod 495:145038. https://doi.org/1016/j.jclepro.2025.145038
66. Sutton RS, Barto AG (2018) Reinforcement learning, second edition: an introduction. MIT Press, Cambridge https://books.google.fi/books?id=sWV0DwAAQBAJ
67. Tul Q, Ali M, Riaz A, Noureen A, Kamranz M, Hayat B, Rehman A (2017) Sentiment analysis using deep learning techniques: a review. Int J Adv Comput Sci Appl 8(6). https://doi.org/14569/IJACSA.2017.080657
68. van Engelen JE, Hoos HH (2020) A survey on semi-supervised learning. Mach Learn 109:373–440
69. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I (2017) Attention is all you need. In: Guyon I, Von Luxburg U, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R (eds) Advances in neural information processing systems, vol. 30. Curran Associates. https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf
70. Vayansky I, Kumar SAP (2020) A review of topic modeling methods. Inform Syst 94:101582. https://doi.org/1016/j.is.2020.101582
71. Varoquaux G, Colliot O (2023) Evaluating machine learning models and their diagnostic value. In: Colliot O (eds) Machine learning for brain disorders. Springer Humana Press, New York, pp 601–630
72. Yang Z, Dai Z, Yang Y, Carbonell J, Salakhutdinov RR, Le QW (2019) XLNet: generalized autoregressive pretraining for language understanding. In: 33rd international conference on neural information processing systems, pp 5754–5764

73. Young DS (2017) Handbook of regression methods. Taylor & Francis, New York
74. Yue L, Chen W, Li X, Zuo W, Yin M (2019) A survey of sentiment analysis in social media. Knowl Inform Syst 60(2):617–663. https://doi.org/1007/s10115-018-1236-4
75. Yu T, Zhu H (2020) Hyper-parameter optimization: a review of algorithms and applications. arXiv preprint arXiv:2003.05689
76. Zhou M, Duan N, Liu S, Shum HY (2020). Progress in neural NLP: modeling, learning, and reasoning. Engineering 6(3):275–290
77. Zhao WX, Jiang J, Weng J, He J, Lim EP, Yan H, Li X (2011) Comparing Twitter and traditional media using topic models. In: Clough P, Foley C, Gurrin C, Jones G J F, Kraaij W, Lee H, Mudoch V (eds) Advances in information retrieval. Springer, Berlin, Heidelberg, pp 338–349
78. Zhang H (2004) The optimality of Naive Bayes. The Florida AI Research Society. https://api.semanticscholar.org/CorpusID:8891634
79. Zhu X (2008) Semi-supervised learning literature survey. Technical Report. 1530, University of Wisconsin-Madison

---

# Footnotes

1 https://github.com/aesuli/SentiWordNet

2 https://github.com/sloria/TextBlob

3 https://github.com/vaderSentiment/vaderSentiment

# Part II
# Framework Design and Numerical Experiments

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2025
M. Švaňa et al., *From Text to Understanding*, Fuzzy Management Methods
https://doi.org/10.1007/978-3-032-00129-0_4

# 4. Framework Design

Miloš Švaňa[1] ✉, František Zapletal[1] and Miroslav Hudec[1]
(1) VSB – Technical University of Ostrava, Ostrava-Poruba, Czech Republic

## 4.1 Introduction to Framework Design

The conducted research addresses the problem of incorporating citizen opinion in municipal decision-making and the problem of dealing with uncertainty associated with these opinions by applying a methodology called Design Science Research (DSR), as mentioned in the introduction to this book. Five stages of DSR from Awareness of Problem to Conclusion are depicted in Fig. 4.1.

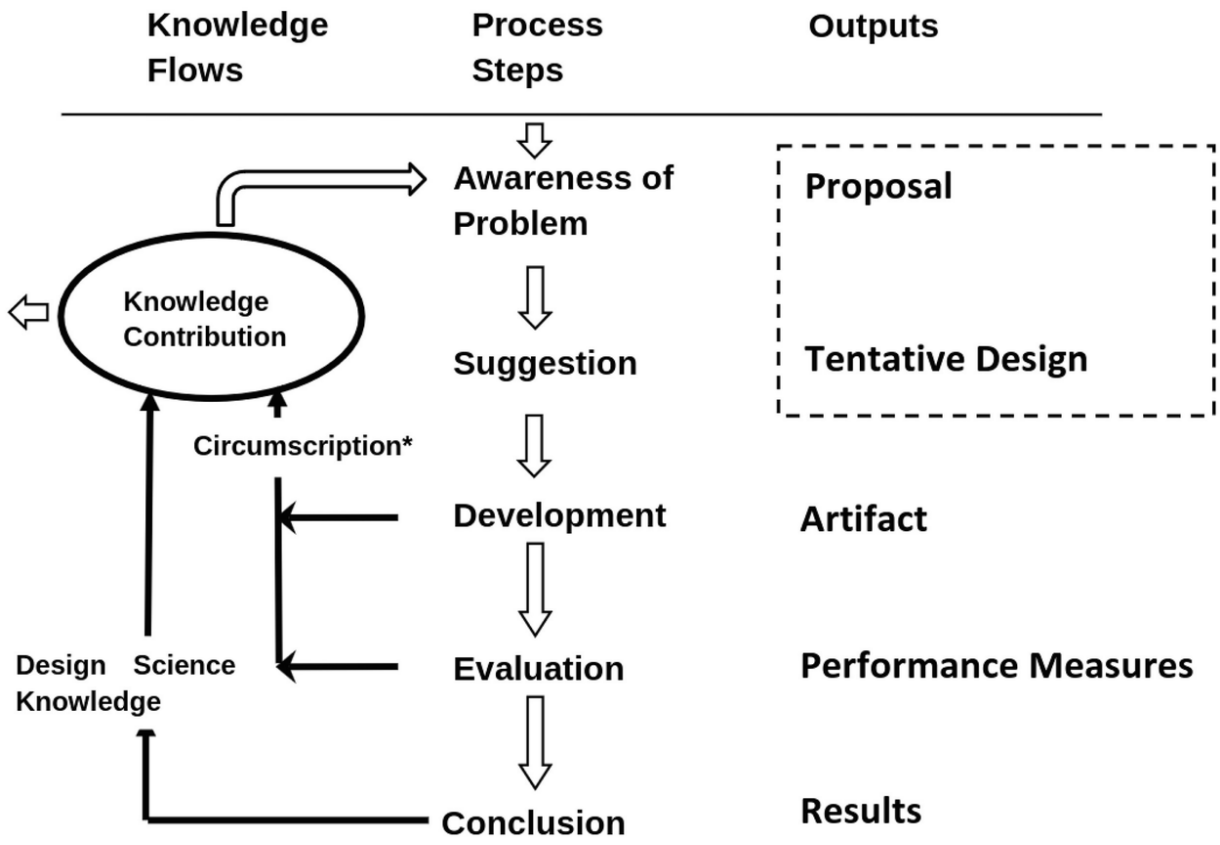


***Fig. 4.1*** Design Science Research process [2]

As declared earlier, the purpose of this work is to propose a framework that would allow municipalities to use social media data as a decision support tool. The framework should be able to identify topics discussed on social networks in a given region and aggregate sentiment expressed toward each topic. This aggregation should capture not only the average sentiment but also its diversity. To fulfill these requirements, the framework includes the following components:

1. A **sentiment analysis** component for determining the overall sentiment polarity of each post

2. A **topic modeling** component to identify topics being discussed in a given municipality and to map individual posts to topics

3. A **fuzzy aggregation** procedure to summarize sentiment expressed toward each identified topic while capturing opinion diversity

4. A calculation of the **level of positive and negative opinion** expressed toward each topic.

Figure 4.2 shows the flow of data between the components of the framework. It also includes all parameters configurable by the framework user (see the three dashed boxes). This high-level view can be thought of as the tentative design produced by the suggestion phase of the DSR process outlined above.

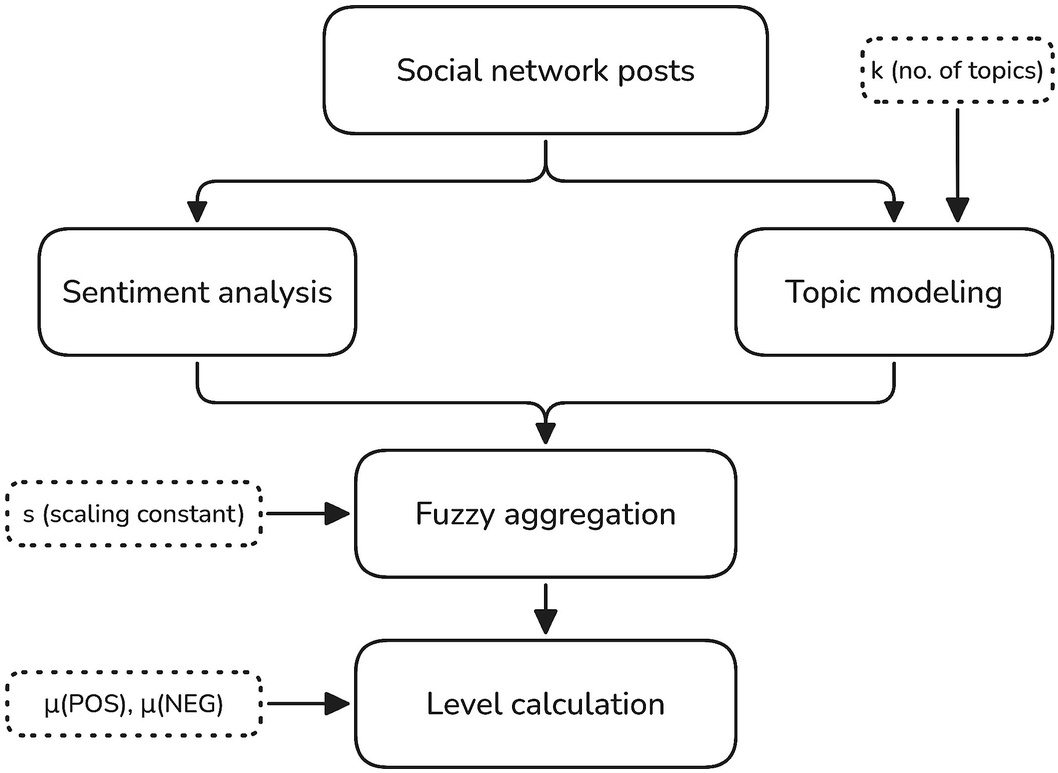


***Fig. 4.2*** Data flow inside the proposed framework (the required inputs from user are presented in dashed boxes)

The remainder of this section justifies the selection of specific sentiment analysis and topic modeling methods. Consequently, it describes fuzzy aggregation and calculation of the level of positive and negative opinions,

which represents the methodological contribution. An explanation of the free parameters in Fig. 4.2 is also provided. Together, the remaining subsections describe the most important aspects of the development phase of the DSR process.

---

## 4.2 Evaluation of Sentiment Analysis Methods

As discussed in Sect. 3.5, three candidates for the role of the sentiment analysis component were considered: (1) lexicon-/rule-based and (2) Naive Bayes sentiment analyzers, both from the TextBlob library and (3) VADER sentiment analysis library. A portion of the Twitter Sentiment Dataset[1] was used to evaluate their performance. This dataset contains tweets related to Indian politics, a topic similar to the intended use case of the proposed framework. Each tweet is labeled as positive (1), negative ($-1$), or neutral (0).

Weighted F1 score[2] was chosen as a performance indicator. Since all three methods produce a continuous output, a discretization procedure had to be applied. For the lexicon-/rule-based TextBlob and VADER analyzers, the overall output polarity ranges from $-1$ (absolutely negative) to 1 (absolutely positive). A simple set of discretization rules is used:

- *IF* polarity $< -0.05$ *THEN* label $= -1$ (negative sentiment)
- *IF* polarity $> 0.05$@ *THEN* label $= 1$ (positive sentiment)
- *OTHERWISE* label $= 0$ (neutral sentiment)

The Naive Bayes analyzer returns the probability that a given document belongs to a positive sentiment class. In order to use the aforementioned rules, the probability of

positive sentiment $P(\mathrm{Pos})$ is transformed into a polarity value *Pol* from a $[-1, 1]$ interval by

$$Pol = 2P(\mathrm{Pos}) - 1. \tag{4.1}$$

In this dataset, the TextBlob lexicon-/rule-based analyzer achieved an F1 score of 0.932. The Naive Bayes and VADER analyzers achieved F1 scores of 0.341 and 0.562, respectively. Therefore, the TextBlob lexicon-/rule-based analyzer was selected to play the role of the sentiment analysis component within the framework. F1 score was also evaluated with the boundaries in the decision rules above set to $-0.1$ and $0.1$. This configuration caused a drop in the F1 score for all analyzers.

Based on these results, the TextBlob lexicon-/rule-based sentiment analyzer was chosen to be incorporated into the framework.

---

## 4.3 Evaluation of Topic Modeling Methods

Two candidate topic modeling methods described in detail in Sect. 3.4 were considered: LDA represents a more traditional approach, while BERTopic employs modern natural language processing and deep learning methods. A dataset of Ostrava tweets was used to evaluate these two options (see further in Sect. 5.2). To ensure the robustness of the results, the evaluation was repeated for different numbers of requested topics ranging from 5 to 47. The number of topics to extract $k$ is one of the free parameters of the framework depicted in Fig. 4.2. As mentioned in Sect. 3.3, the methods were compared on average normalized pointwise mutual information (NPMI) and the average ratio of unique words (topic diversity).

According to the results shown in Fig. 4.3, BERTopic dominates in both measures, the only exception being a slightly lower topic diversity score when the number of

topics was set to 11. Ignoring very low numbers of topics, one can also observe that NPMI seems to decrease as the number of topics increases. The diversity of topics exhibits the opposite trend. Given these results, BERTopic was chosen to be incorporated into the framework prototype.

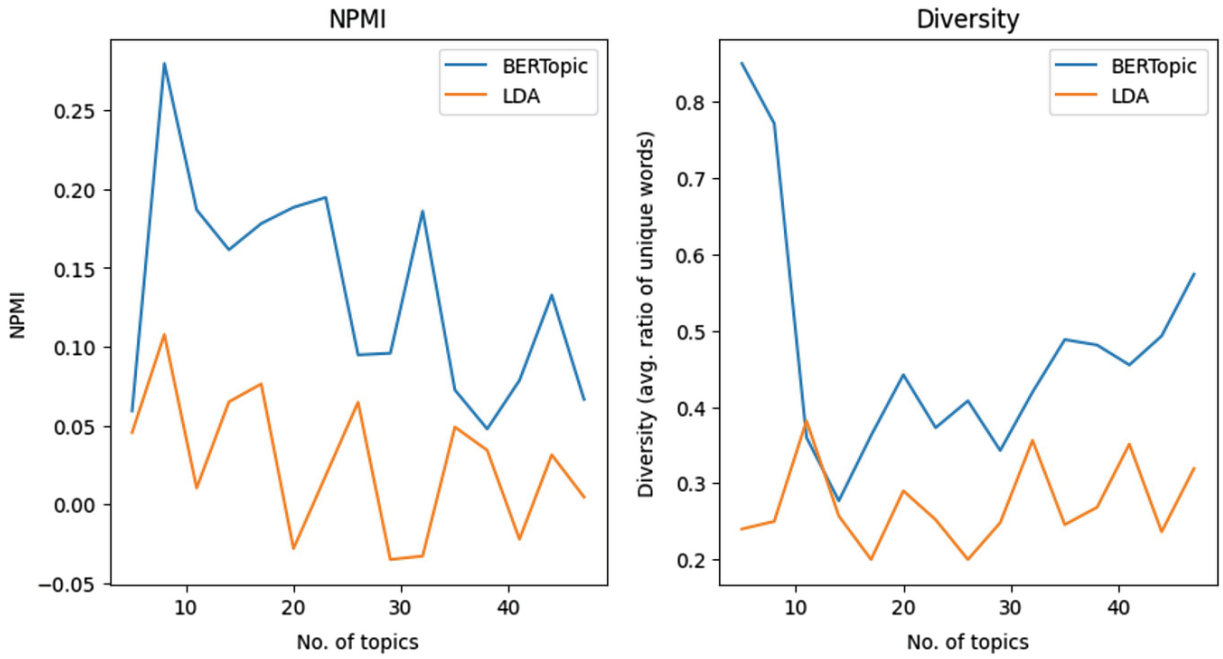


***Fig. 4.3*** Comparison of NPMI and topic diversity between LDA and BERTopic for various numbers of topics

## 4.4 Fuzzy Sentiment Aggregation

Information about tweet sentiment and their assignment to topics must be aggregated in a form that is easy to understand by decision-makers. However, simple methods such as the arithmetic mean lead to a loss of information about sentiment diversity, which might be useful in many scenarios.

To address this issue, the sentiment toward each topic is modeled as a triangular fuzzy number (TFN) (2.3). Assuming that both the sentiment polarity and the topic distribution of each social media post are known, the core of this triangular fuzzy number is calculated as a weighted

mean, with topic probabilities as a support of TNF (Sect. 2.1.1). The polarities and topic distributions from Table 4.1 for three fictional social media posts can be considered as an example.

***Table 4.1*** Example of an input for the aggregation procedure

| | | Topic distribution | | |
|---|---|---|---|---|
| | **Polarity** | **Topic 1** | **Topic 2** | **Topic 3** |
| **Post 1** | 0.5 | 0.5 | 0.3 | 0.2 |
| **Post 2** | 0.35 | 0.3 | 0.4 | 0.3 |
| **Post 3** | −0.2 | 0.2 | 0.2 | 0.6 |

The core of the TFN representing the sentiment toward topic 1 is calculated as

$$m_{t1} = \frac{0.5 \cdot 0.5 + 0.35 \cdot 0.3 - 0.2 \cdot 0.2}{0.5 + 0.3 + 0.2} = 0.315. \tag{4.2}$$

The support of sentiment modeled by TFN should provide a reliable representation of real sentiment diversity. Higher diversity should lead to a wider support for TFN. An intuitive approach would be to use statistical measures such as variance. However, variance cannot capture asymmetry, which can lead to an incorrect representation of reality.

After examining multiple options, the square root of semivariance ("standard semideviation") was selected as a basis for determining the support interval. Weighted semivariance for topic $t$ and side $s \in \{\text{left}, \text{right}\}$ can be calculated as

$$\text{SVar}_{t,s} = \frac{\sum_{i=1}^{N} w_i (x_i - x^*)^2}{\frac{M-1}{M} \sum_{i=1}^{N} w_i}, \tag{4.3}$$

where $N$ represents the total number of points on the left (or on the right) from the mean, $x^*$ represents weighted mean, and $M$ represents the number of nonzero weights. Moreover, all $x_i \leq x^*$ if $s = left$, and all $x_i \geq x^*$ if $s = \text{right}$ .

Once the weighed semivariance is determined for both sides, the support interval of the sentiment TFN can be calculated as

$$[m_t - s\sqrt{\mathrm{SVar}_{t,l}}; m_t + s\sqrt{\mathrm{SVar}_{t,r}}], \tag{4.4}$$

where $s$ is another free parameter of the framework—a positive real number determining the relationship between the "breadth" of the fuzzy number and weighted "standard semideviation." Its influence on the result of the aggregation is the subject of further research. The indices $l$ and $r$ identify the left and right sides of the mean.

Figure 4.4 shows the TFNs constructed for several edge-case sentiment distributions. The maximum "breadth" is achieved in case (e) with all an equal number of samples with polarities $-1$ and $1$. Uniform distributions in cases (a), (b), and (c) lead to a TFN that spreads throughout the sample range. Case (f) shows a simple uneven sentiment distribution. Since there are more samples with sentiment $0.25$ than samples with sentiment $-0.5$, there is less uncertainty on the right side.

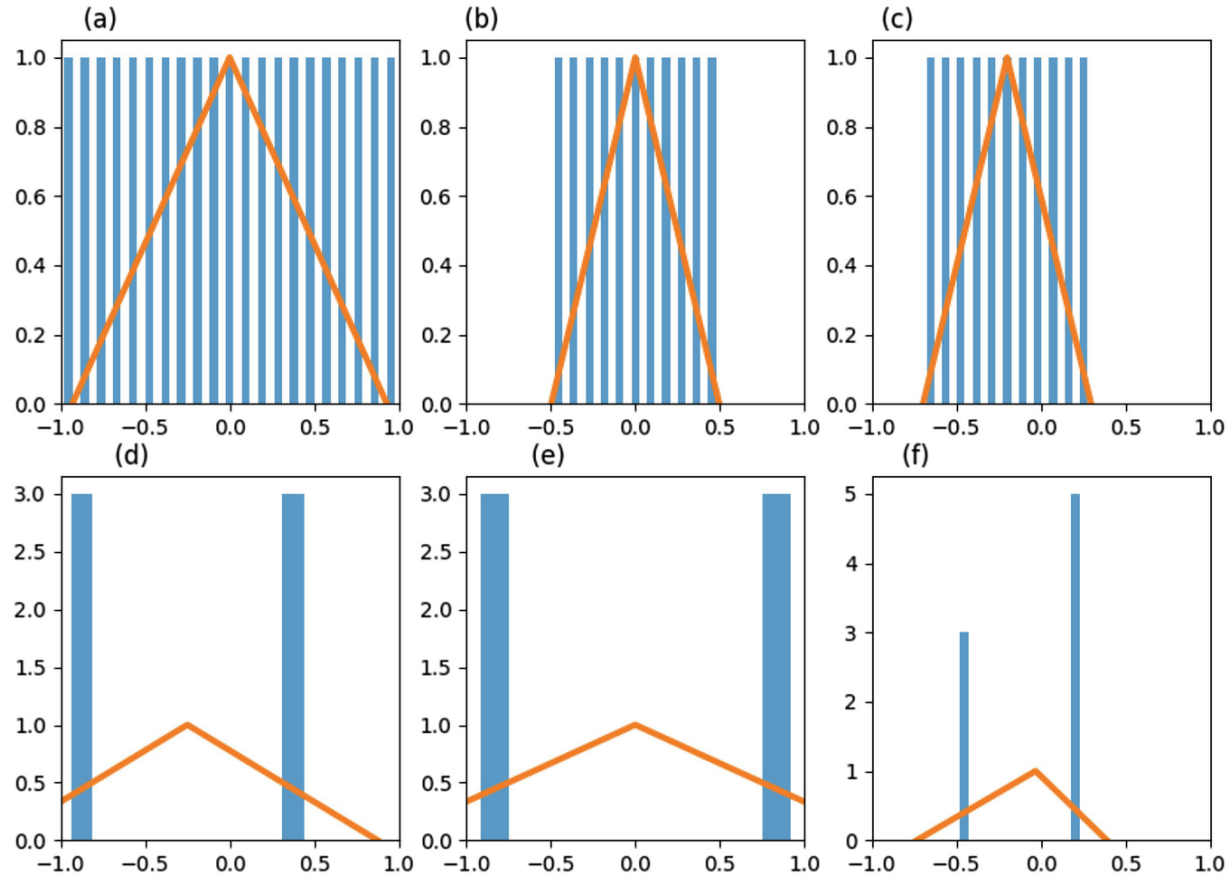


***Fig. 4.4*** TFNs representing sentiment constructed for various edge-case distributions of individual post sentiments ($s = 1.5$). (**a**) Uniform from $-1$ to 1. (**b**) Uniform from $-0.5$ to 0.5. (**c**) Uniform shifted to left. (**d**) Only $-1$ and 0.5. (**e**) Only $-1$ and 1. (**f**) Uneven distribution

---

# 4.5 Levels of Positive and Negative Opinions

The TFN representation can be difficult to understand for a layperson. Therefore, the final step performed by the framework is the calculation of the level of positive and negative opinion expressed toward a given topic. One possible approach to determine these levels is to calculate them as degrees of similarity between the topic TFN and fuzzy sets representing the concepts of positive and negative opinions.

The shapes of the fuzzy sets of positive and negative opinions are the last free parameters of the model. To carry out the case studies presented in this dissertation, the two shapes were estimated from a small-scale survey, in which respondents determined the level of positive and negative opinions on randomly selected tweets. All tweets were in Czech, and all respondents were Czech speakers. In total, 420 tweet evaluations were collected.

TextBlob sentiment polarity was used as the universe of discourse, which matches the universe of discourse of the topic TFNs presented in Sect. 4.4.

The mean squared error minimization was used to estimate the shapes of fuzzy sets. It was assumed that the membership function of the positive opinion fuzzy set starts at zero at polarity $a$, grows linearly, reaches the value of 1 at polarity $b$, and then stays constant until polarity 1.0. The membership function of the negative opinion fuzzy set mirrors this shape. This shape assumption significantly simplifies the calculation of similarity.

The estimated shapes are shown in Fig. 4.5. They are not symmetric; the negative opinion fuzzy set reaches 0 at polarity 0.75, but the positive opinion fuzzy set starts increasing at polarity $-0.32$. It could be speculated that respondents are more sensitive to negative aspects of evaluated tweets.

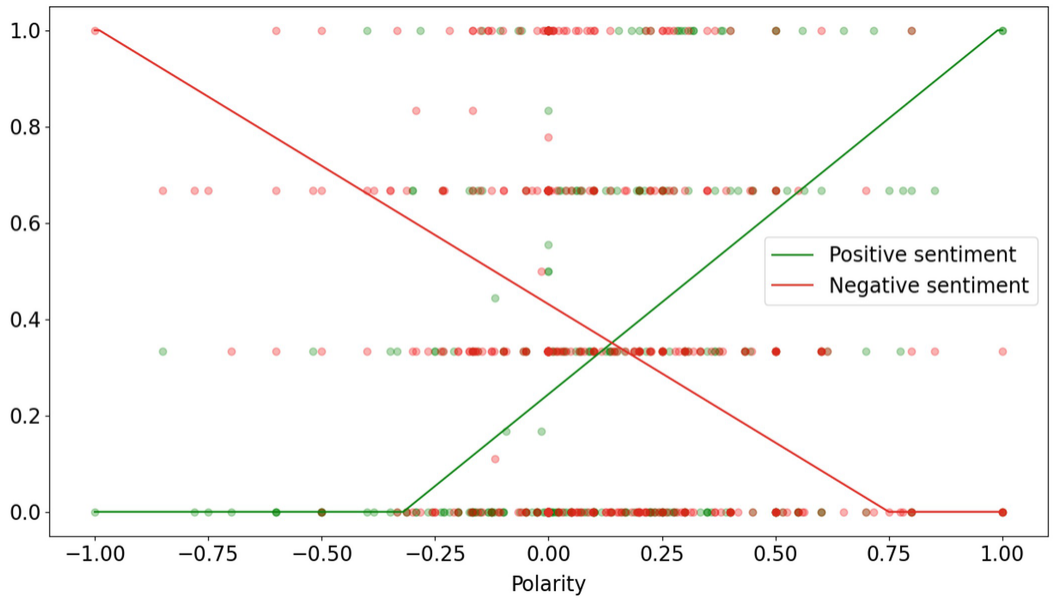


***Fig. 4.5*** Fuzzy sets representing positive (green) and negative (red) opinion derived from survey data

The similarity between fuzzy and crisp sets can be measured in different ways. Jaccard index is a common measure of crisp sets similarity and could be fuzzyfied by using some variant of fuzzy union, intersection, and cardinality. However, the Jaccard index would be positively correlated with opinion diversity, which is undesired in the presented use case.

Therefore, a measure of similarity, already shown in Sect. 2.2 (2.47), is used [1].

The equation calculates the maximal degree of membership of the standard intersection of the two fuzzy sets. This measure has several desirable properties: It is commutative, it is bounded by the interval $[0, 1]$, and it reaches the maximum value of 1 if $\tilde{A} = \widetilde{B}$ and the minimum value of 0 if $\tilde{A} \cap \widetilde{B} = \emptyset$. Given the shapes of positive and negative opinion sets, the similarity value increases only if tweets with appropriate sentiment polarity are added to the topic. Moreover, the measure can also be interpreted as the possibility of the topic sentiment TFN being “less than negative opinion” or being “greater than positive opinion.”

It is also worth mentioning that the measure was also used by [3] in the evaluation of the opinions of stakeholders in education during the different phases of COVID, where hesitance in answers was handled with the support of TFN.

# References


1. Lee-Kwang H, Song YS, Lee KM (1994) Similarity measure between fuzzy sets and between elements. Fuzzy Sets Syst 62(3):291–293 [Crossref]
2. Vaishnavi V, Kuechler B, Petter S (2004) Design science research in information systems. http://www.desrist.org/design-research-in-information-systems/
3. Zapletal F, Hudec M, Švaňa M, Němec R (2023) Three-level model for opinion aggregation under hesitance. Soft Comput. https://doi.org/1007/s00500-023-07853-2


---

# Footnotes

1 https://www.kaggle.com/datasets/saurabhshahane/twitter-sentiment-dataset

2 F1 score is a harmonic mean of precision and recall. Precision is the ratio of true positives to all items labeled positive. Recall is the ratio of true positives to all real positives.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2025
M. Švaňa et al., *From Text to Understanding*, Fuzzy Management Methods
https://doi.org/10.1007/978-3-032-00129-0_5

# 5. Numerical Experiments

Miloš Švaňa[1] ✉, František Zapletal[1] and Miroslav Hudec[1]
(1) VSB – Technical University of Ostrava, Ostrava-Poruba, Czech Republic

## 5.1 Overview

This chapter corresponds with the evaluation step of the DSR process. To the authors' knowledge, there are no standardized benchmarks for the task the framework is attempting to accomplish. Therefore, case studies and basic comparative analyses were chosen as the main evaluation methods.

Table 5.1 provides an overview of all experiments performed with the framework. First, the framework was applied to a synthetic dataset. The goal was to verify that the results produced by the framework match the expectations derived from the known structure of the data in this dataset. Second, the framework was applied to real data from Ostrava and Brno to analyze what useful information can be extracted and how the two cities compare. Third, the influence of including likes and retweets on the resulting TFNs was analyzed. Fourth, the framework was compared with a naive method, and it is shown that it provides more information. Finally, this chapter introduces a prototype of a simple web application

to present the results provided by the framework in a user-friendly way.

***Table 5.1*** List of experiments performed to evaluate the framework

| Main question | Type | Datasets |
|---|---|---|
| **Does the framework extract expected information from synthetic data?** | Case study | Synthetic |
| **What interesting information can the framework extract from real-world data?** | Case study | Ostrava and Brno |
| **Can the framework results be used to compare two cities?** | Case study | Ostrava and Brno |
| **What is the influence of likes and retweets on the results?** | Comparative analysis | Ostrava |
| **Does the framework provide more information than a simpler method?** | Comparative analysis | Ostrava |

### 5.1.1 Synthetic Dataset

The synthetic dataset was generated with ChatGPT. It was asked multiple times to generate a social media post in Czech. The prompt specified the topic of the tweet and sentiment. The dataset contains:

- Six very positive posts on healthcare.
- Three slightly positive posts on healthcare.
- Three very negative posts on healthcare.
- Three very negative posts on education.
- Two slightly negative posts on education.

Given this distribution, the framework should show that compared to education, sentiment toward healthcare is more diverse and also more positive on average.

Starting with sentiment analysis, the framework was able to correctly evaluate sentiment polarity, at least in a relatively short way. As shown in Table 5.2, very positive synthetic tweets have the highest average sentiment polarity. They are followed by slightly positive, slightly negative, and finally very negative tweets. However, one

can also observe that the average polarity is biased toward positive values, as even very negative tweets have an average polarity of $-0.135$.

***Table 5.2*** Average polarities of different topics in the synthetic dataset

| Known sentiment label | Average sentiment |
|---|---|
| **Very negative** | -0.135 |
| **Slightly negative** | 0.101 |
| **Slightly positive** | 0.144 |
| **Very positive** | 0.496 |

The groups are not equidistant. The difference in average sentiment is relatively low between slightly positive and slightly negative tweets, but relatively high between slightly negative and very negative tweets, as well as between slightly positive and very positive tweets. It is possible that the language model used to generate the synthetic dataset interprets the word “slightly” as very close to neutral.

The framework correctly detected expected topics and labeled them as *education, outdated, czech* and *healthcare, czech, czech republic*. References to the Czech Republic make sense since the language model was tasked with generating tweets about healthcare and education in the Czech Republic. Similarly, the use of the word *outdated* to describe the topic of education is not surprising, because the dataset purposefully contains only negative tweets on the topic.

The probabilities of different tweets that belong to the detected topics also match the expectations. The probabilities of belonging to the topic of healthcare range from 0.46 to 1.0 for tweets that actually belong to this topic and from 0.0 to 0.12 for tweets that do not. The situation is similar for the topic of education with membership probabilities ranging from 0.86 to 1 for tweets that actually

belong to the topic and from 0.0 to 0.25 for tweets that do not belong to the topic. The sum of probabilities for each individual tweet can be less than 1, as there can be a nonzero probability of the tweet being an outlier and not belonging to either of the two topics. There are also tweets with relatively high probabilities of belonging to both topics. Since there is some semantic overlap between the topics (both topics are related to the Czech Republic), this result is expected.

In the next step, the results of sentiment analysis and topic modeling were used to construct a TFN for each topic. These are depicted in Fig. 5.1.

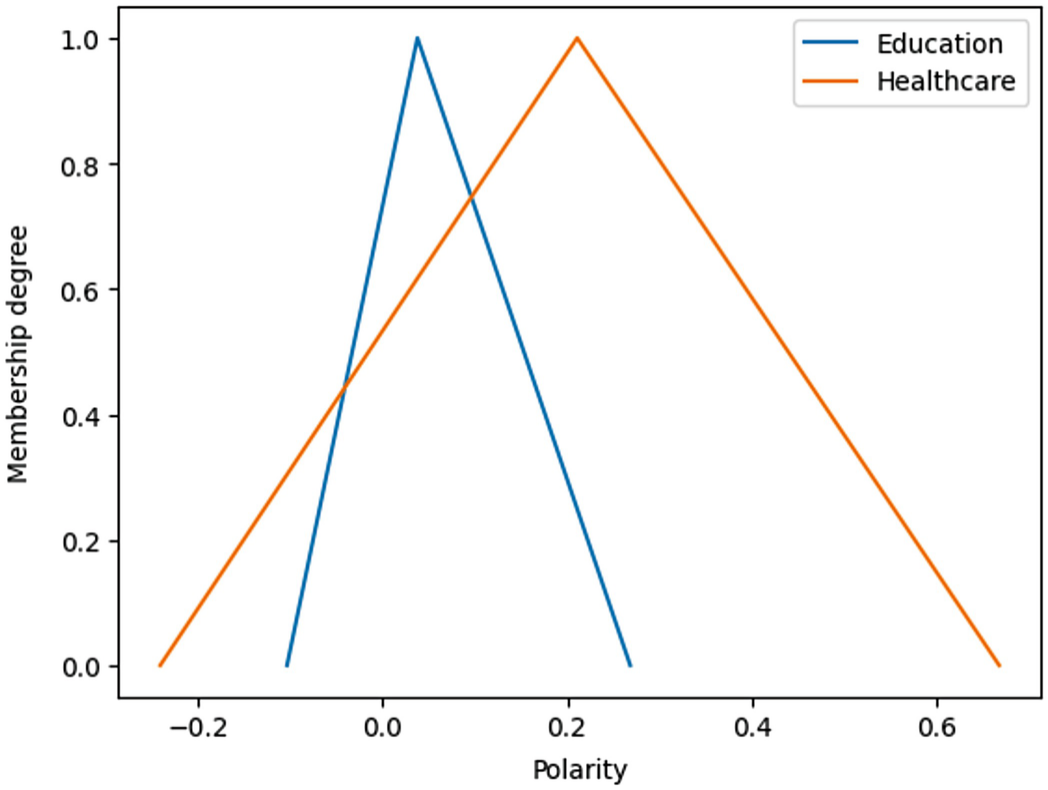


***Fig. 5.1*** Fuzzy sets representing the aggregated sentiment toward the two topics in the synthetic dataset

As expected, the mean sentiment represented by the cores of the two TFNs is lower for education, because all tweets related to this topic are very or slightly negative.

The support of the TFN representing healthcare is much wider, as the input dataset contains a mixture of positive and negative tweets related to healthcare. It is also possible to observe that the healthcare TFN is almost symmetric. In contrast, the education TFN is more asymmetric, with the core shifted to the right. The asymmetry reflects a larger number of tweets with sentiment polarity greater than the mean.

Finally, TFNs were used to determine the level of positive and negative sentiment expressed toward the two topics. The levels of positive and negative sentiment are 0.379 and 0.451, respectively, for education and 0.491 and 0.359, respectively, for healthcare. This result is compatible with expectations and with the shapes of the TFNs. However, the level of positive sentiment in education is surprisingly high, given that all tweets about this topic are negatively oriented. Some slightly negative tweets mention several positive aspects, which might at least partly explain this result.

---

## 5.2 Case Study of Ostrava and Brno

The functionality of the framework is demonstrated on data obtained from Twitter (now also called X). Twitter is a popular social media platform around the world. Users of this network interact by publishing short pieces of content often called *tweets*. Tweets can contain text, images, videos, links, or even simple surveys. By default, all tweets are public and accessible to any other Twitter user. Twitter users extensively utilize hashtags—labels starting with a # symbol that can be used to categorize and explore tweets. Users can also engage by *liking*, *sharing*, or *responding* to tweets created by others.

Although social networks such as Facebook and Instagram can be assumed to be more popular, Twitter was

chosen mainly because of its well-designed application programming interface (API), which allows researchers and developers to programmatically access its data. In the past, this access was free for research purposes, but in 2023 Twitter removed this option, and at the time of writing, even academic institutions have to pay an API usage fee.

Datasets containing tweets from two Czech cities were used to demonstrate the functionality of the proposed framework. The first dataset contains 19,764 tweets published from Ostrava (population around 280,000) between January 22 and March 17, 2023. The second dataset contains 22,767 tweets published from Brno between January 20 and March 30, 2023 (population around 400,000). The tweets were downloaded from the official Twitter API using search queries `place:"Ostrava" -is:retweet` and `place:"Brno" -is:retweet`. The dataset contains five columns (attributes): text of the tweet, unique author ID, time of publication, number of likes, and number of retweets. The characteristics of dataset match the intended use of the framework as a decision support tool for municipal decision-makers.

In addition to the two real-world Twitter datasets, a small synthetic dataset was created to verify the basic idea of the proposed framework. This dataset contains ten tweets generated by OpenAI's ChatGPT. Five of these tweets belong to a topic labeled *education*, and five tweets belong to a topic labeled *healthcare*. ChatGPT was tasked with generating tweets with different levels of sentiment described by phrases such as *very positive* or *slightly negative*. Given that topics and levels of sentiment are known beforehand, this small dataset can be used to examine whether the framework provides expected results.

In principle, the framework can be used to process any dataset that contains sentiment-colored documents created by multiple users. Examples include data from other social

networks, such as Reddit, open-ended questions from surveys, or product reviews.

### 5.2.1 Preprocessing

Given the location, the vast majority of tweets in the dataset are written in Czech. However, other languages such as English, Slovak, Polish, or Ukrainian are also present. Mixing multiple languages could cause problems during analysis. Moreover, most of the sentiment analysis and topic modeling tools are designed for English. Therefore, a state-of-the-art automatic translation tool DeepL[1] was used to unify the language by translating all Tweets into English (Sect. 3.2.3).

Translating can lead to some loss of information. For example, DeepL is not capable of dealing with words from local dialects. The following tweet contains two such words in Czech: *bo* (meaning: because) and *sichtu* (meaning: work shift):

> Bo ochutnala moc vína a jiných dobrot z Jižní Moravy.
> A dokonce jela moc brzo ráno zpět na sichtu - do školy.

DeepL translates these two words:

> Bo tasted a lot of wine and other goodies from South Moravia.
> And she even went back to sichta - to school - very early in the morning.

An acceptable translation could look like this:

> Because she tasted many wines and other goodies from South Moravia.
> And she even went back to her shift - to school - very early in the morning.

Since TextBlob analyzes patterns, such as the use of capital letters, there was no additional preprocessing for sentiment analysis. However, a sequence of standard text preprocessing procedures was applied before BERTopic was used to detect topics and assign individual tweets to these detected topics. The preprocessing sequence includes the following steps [3]:

- Turning the whole text into lowercase
- Removal of unnecessary whitespace characters, e.g., new line characters or sequences of multiple spaces
- Removal of all non-alphanumeric characters, including the hashtag symbol or punctuation
- Expansion of contractions, e.g., *don't* was replaced with *do not*
- Removal of stopwords, i.e., commonly used words carrying little to no meaning
- Lemmatization, i.e., transformation into a basic dictionary form

### 5.2.2 Basic Data Analysis

This section covers a basic analysis of data related to real-world tweet datasets from Brno and Ostrava. In this dataset, six quantitative variables are examined: tweet sentiment, absolute tweet sentiment (sentiment intensity disregarding its orientation), number of likes, number of retweets, post-length (number of characters), and number of posts per author.

Figure 5.2 shows the histograms for all variables with the exception of absolute tweet sentiment for the Ostrava dataset. The tweet sentiment somewhat approximates a normal distribution. Most of the tweets in the dataset seem to have a sentiment value close to 0; they are neutral. Other variables seem to approximate an exponential distribution. Most of the posts have a low number of likes and retweets and are relatively short. Similarly, most

authors publish a relatively low number of posts. However, it is possible to notice a sudden spike at the right end of the post-length histogram. This spike is most likely caused by the post-length limit imposed by Twitter (280 characters). It is possible that a large number of people wanted to create a longer post but were forced to make it shorter to stay within the post-length limit.

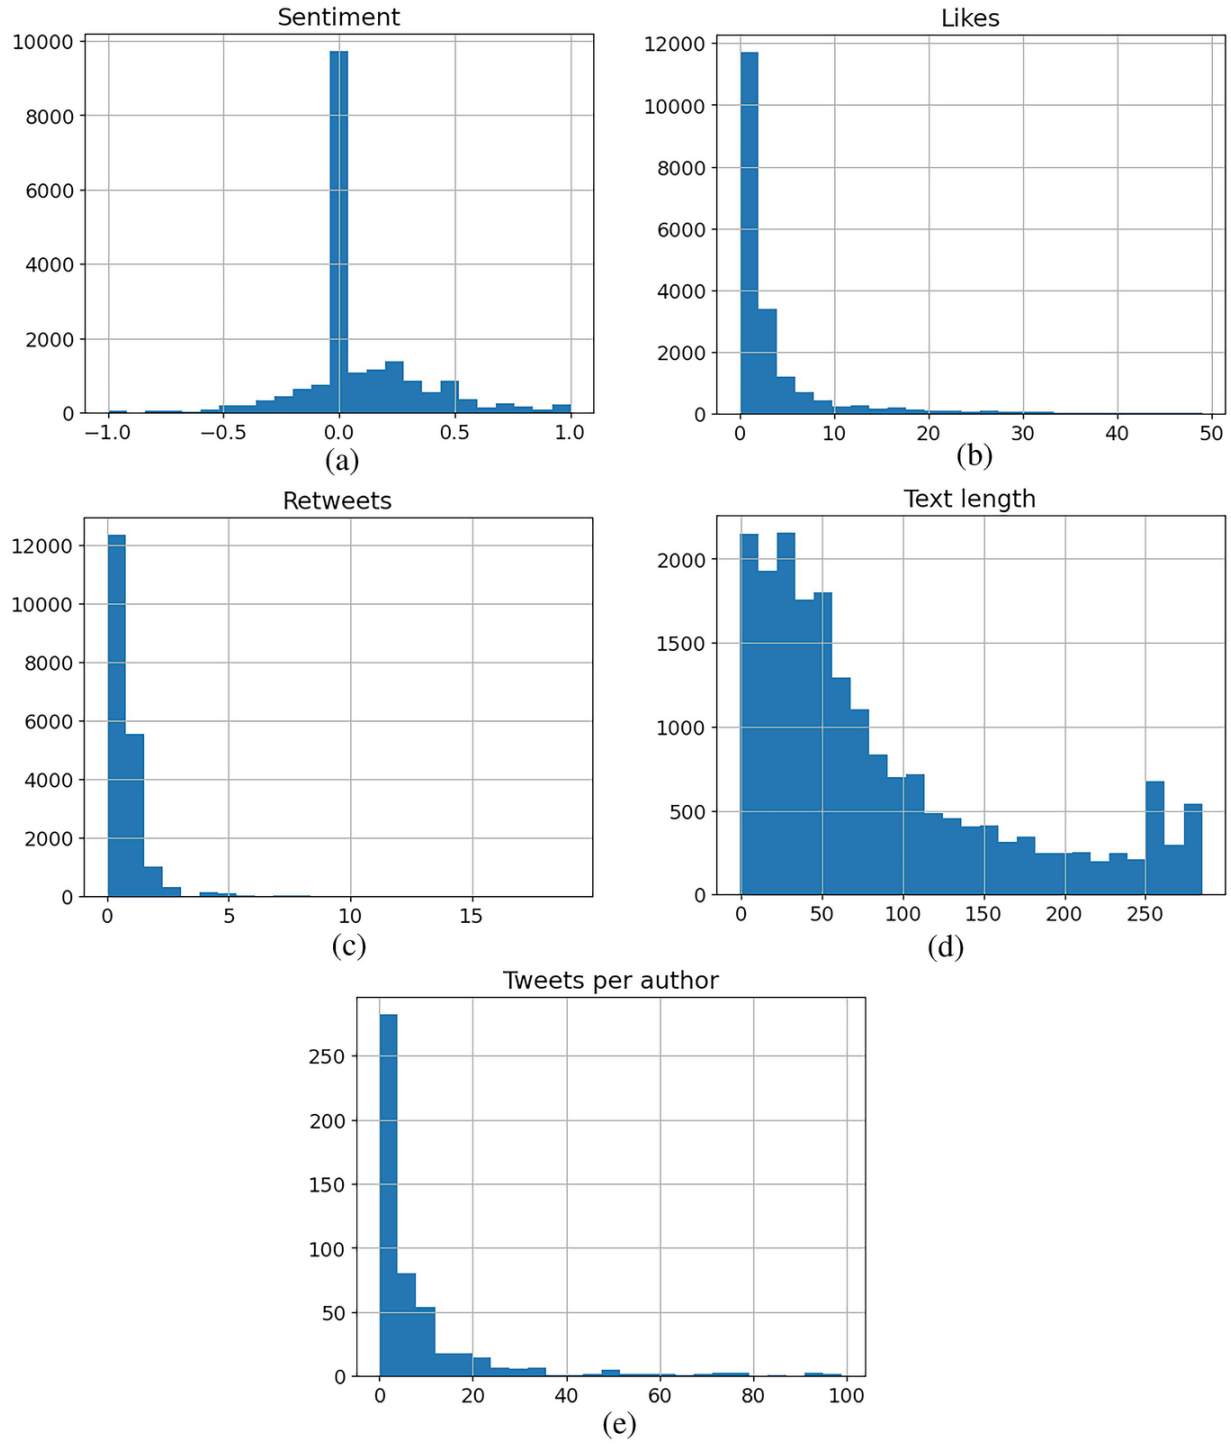


***Fig. 5.2*** Histograms for variables related to tweets from Ostrava. Note: The histograms do not show several extreme values. (**a**) Sentiment. (**b**) Likes. (**c**) Retweets. (**d**) Post length. (**e**) Posts per author

Table 5.3 lists several basic descriptive statistics for the same five variables. The average sentiment is slightly greater than 0, but the median is 0. This suggests that the

data are skewed to the left. It can also be observed that Twitter users seem to use the like functionality significantly more than the retweet functionality. At the same time, the positive mean value for sentiment is rather unexpected, as many people anecdotally associate social networks with negative sentiment.

***Table 5.3*** Descriptive statistics of variables related to tweets from Ostrava

| | Sentiment | Likes | Retweets | Text length | Posts per author |
|---|---|---|---|---|---|
| **Mean** | 0.086 | 8.10 | 0.74 | 85.48 | 35 |
| **Std. dev.** | 0.267 | 80.73 | 4.56 | 78.50 | 169 |
| **Minimum** | -1.0 | 0 | 0 | 0 | 1 |
| **Median** | 0.0 | 1 | 0 | 57 | 3 |
| **Maximum** | 1.0 | 4171 | 304 | 285 | 2956 |

Half of the posts received only 0–1 likes and no retweets. The most notable information about text length is that there are posts with a length of 0. This can be explained by the possibility of posting multimedia content, such as images or videos, without including any textual description. Finally, it can be observed that there is a single user who published almost 3,000 posts. This number represents about 15% of the whole dataset. One would expect that this user account would belong to a news agency or a municipal organization. However, the tweets associated with a given author ID suggest that the account belongs to a very active person.

It is also beneficial to investigate the correlation between the examined variables. Table 5.4 shows the Pearson correlation between five variables. The number of posts per author is excluded because it does not make sense to investigate the correlation between variable values for individual tweets and variable values for data aggregated according to the author ID dimension. On the other hand, the table includes absolute sentiment

describing sentiment intensity without disregarding its polarity.

***Table 5.4*** Correlation variables related to tweets from Ostrava

| | Sentiment | Absolute sentiment | Likes | Retweets | Post-length |
|---|---|---|---|---|---|
| **Sentiment** | | ***0.512 | 0.0086 | −0.0121 | 0.0131 |
| **Absolute sentiment** | | | 0.0112 | *0.0177 | ***0.074 |
| **Likes** | | | | ***0.709 | ***0.0445 |
| **Retweets** | | | | | ***0.0685 |

Interestingly, there is a rather strong statistically significant correlation between sentiment and absolute sentiment, as well as between the number of likes and retweets. It is reasonable to assume the presence of a positive feedback loop between the latter pair of variables. In contrast, a statistically significant positive correlation between post-length and absolute sentiment, as well as number of likes and number of retweets, is rather unexpected. At this point, it is not possible to offer a reasonable explanation for this correlation. Finally, there is a statistically significant but rather weak correlation between the number of retweets and the absolute sentiment. It is reasonable to expect emotionally intense content to be retweeted more than neutral content.

Figure 5.3 and Table 5.5 provide a basic description of the variables related to the dataset of tweets published from Brno. Histograms suggest that the variables are distributed similarly to their Ostrava counterparts. However, the descriptive statistics in Table 5.5 reveal some differences between the two cities. The mean sentiment in Brno, as well as the mean number of likes and retweets, is slightly lower. At the same time, the standard deviation and maximum values for these three variables suggest higher diversity. Both median and mean values of text length

suggest that people in Brno publish shorter posts. Individual users also publish fewer posts.

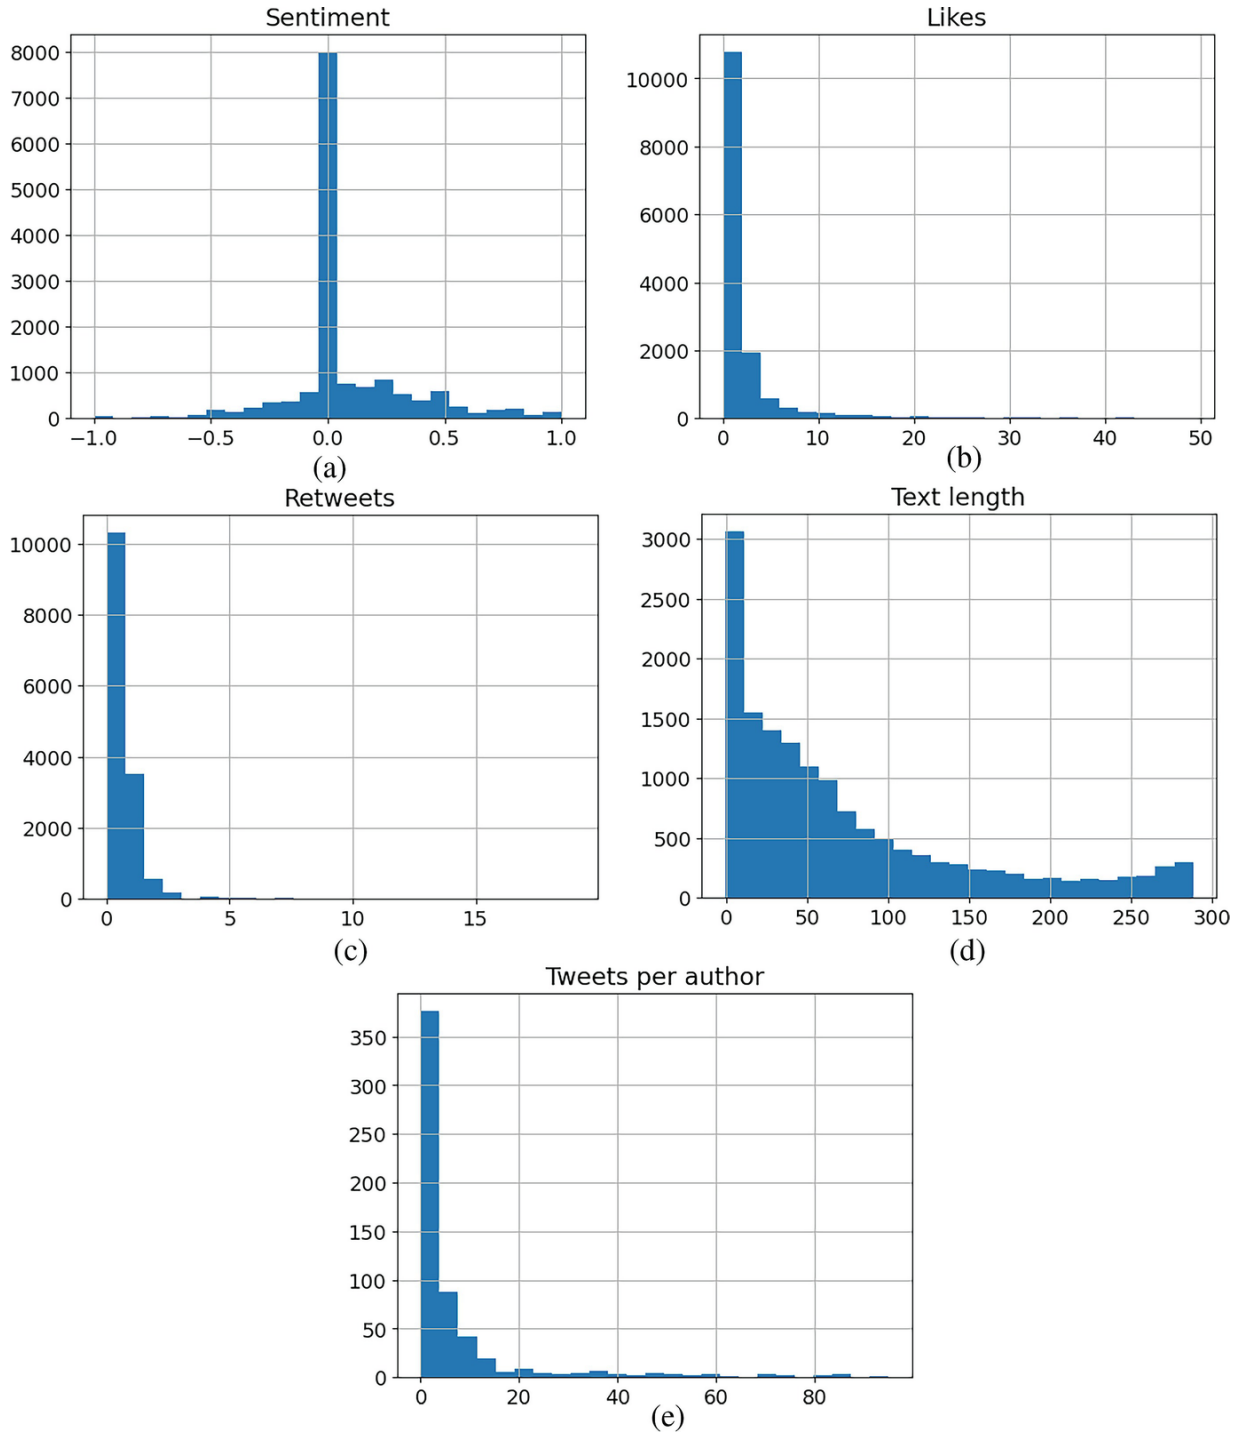


***Fig. 5.3*** Histograms for variables related to tweets from Brno. (**a**) Sentiment. (**b**) Likes. (**c**) Retweets. (**d**) Post length. (**e**) Posts per author

***Table 5.5*** Descriptive statistics of variables related to tweets from Brno

| | Sentiment | Likes | Retweets | Text length | Posts per author |
|---|---|---|---|---|---|
| **Mean** | 0.076 | 6.08 | 0.63 | 72.67 | 23.55 |
| **Std. dev.** | 0.267 | 85.72 | 4.88 | 76.38 | 103 |
| **Minimum** | -1.0 | 0 | 0 | 0 | 1 |
| **Median** | 0.0 | 1 | 0 | 47 | 2 |
| **Maximum** | 1.0 | 7841 | 382 | 288 | 1847 |

Finally, Table 5.6 shows a Person correlation matrix for Brno data. It is again possible to observe several significant differences between the two cities. The correlation between the number of likes and retweets is much weaker in Brno. However, the correlation between post-length and absolute sentiment, as well as retweets, is stronger. Moreover, the correlation between post-length and (originally polarized) sentiment is now statistically significant. At the same time, the statistical significance between absolute sentiment and the number of retweets has been lost.

***Table 5.6*** Correlation variables related to tweets from Brno

| | Sentiment | Absolute sentiment | Likes | Retweets | Post-length |
|---|---|---|---|---|---|
| **Sentiment** | | ***0.491 | 0.0022 | 0.0033 | ***0.036 |
| **Absolute sentiment** | | | 0.0049 | 0.0133 | ***0.145 |
| **Likes** | | | | ***0.310 | ***0.0353 |
| **Retweets** | | | | | ***0.0791 |

Confirming the statistical significance of the differences between the two cities and providing an explanation are beyond the scope of this book. However, these differences could serve as a trigger for future sociological research.

### 5.2.3 Tweets from Ostrava

In the real-world dataset of Ostrava tweets, the framework identified 210 different topics (the number of topics to return was left unspecified when applying the framework). Figure 5.4 shows the TFNs representing the aggregated sentiment toward the eight most discussed topics, as well as the fuzzy sets representing the positive and negative opinions.

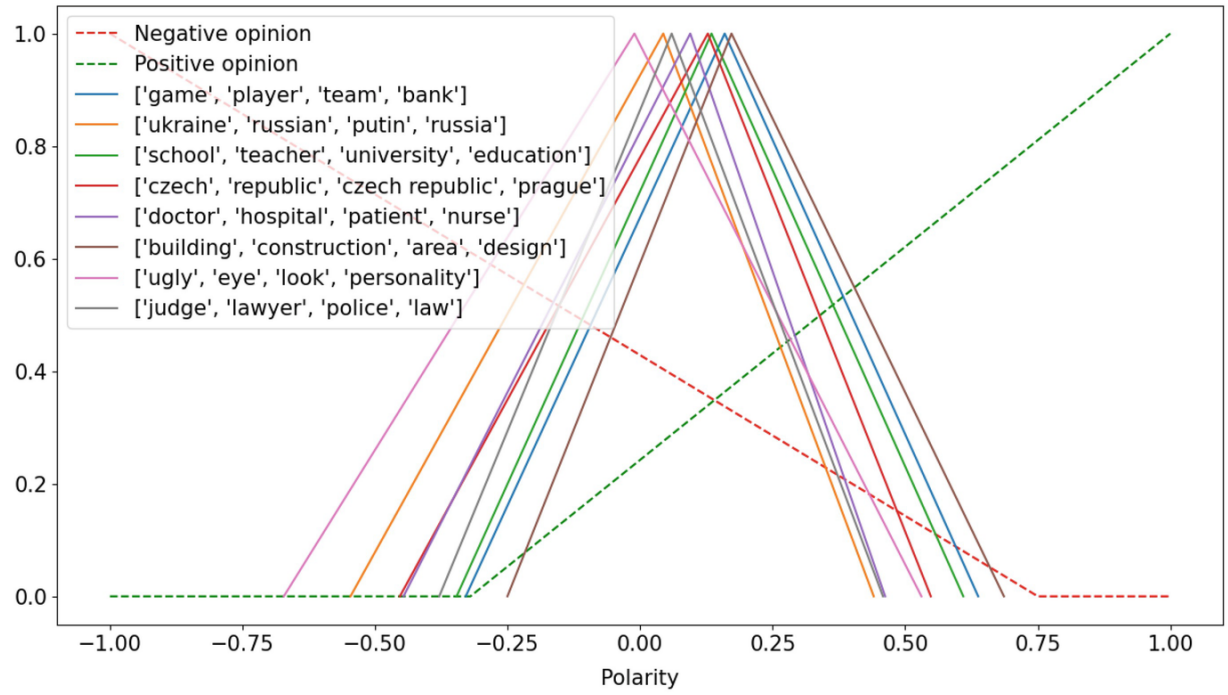


***Fig. 5.4*** TFNs representing aggregated sentiment toward selected topics detected in tweets published from Ostrava ($s = 1.5$)

The topics *school, teacher, university, education* and *doctors, hospital, patient, nurse* are important for municipal decision-making. The shapes of their TFNs suggest that the former is perceived more positively. The diversity of opinions on both topics seems to be quite similar. When examining the exact values, one can conclude that the support for the first topic is slightly wider (0.955) than the support of the second topic (0.908); the first topic is slightly more controversial.

The topic perceived most positively is *building, construction, area, design*, and the topic perceived most negatively is *ugly, eye, look, personality*. The latter topic also seems to be the most controversial among the top eight topics displayed in the figure.

Figure 5.4 can be difficult to read for someone lacking the necessary background knowledge. Table 5.7 shows an alternative representation: the level of positive and negative opinion scaled to a unit interval. It compares several topics that are potentially relevant to municipal decision-makers. The data suggest that people have a relatively positive opinion on the construction of new buildings, reading books, or activities related to elderly citizens. On the other hand, cars and parking, as well as healthcare, seem to be perceived rather negatively and should be prioritized by decision-makers.

***Table 5.7*** Mean sentiment and scaled level of positive and negative opinion for several topics relevant to municipal decision-making in Ostrava

| Topic | Mean sentiment | Level positive | Level negative |
|---|---|---|---|
| **School, teacher, university** | 0.135 | 0.447 | 0.315 |
| **Doctor, hospital, patient** | 0.094 | 0.175 | 0.468 |
| **Building, construction, area** | 0.172 | 0.602 | 0.160 |
| **Car, parking, pedestrian** | 0.069 | 0.186 | 0.567 |
| **Hydrogen, technology, h2** | 0.136 | 0.369 | 0.422 |
| **Book, read, movie** | 0.169 | 0.486 | 0.226 |
| **Old, age, senior** | 0.177 | 0.439 | 0.159 |
| **Poruba, tram, line** | 0.140 | 0.314 | 0.248 |
| **Planetarium, transformation, night** | 0.121 | 0.440 | 0.389 |
| **Snow, snowing, winter** | 0.142 | 0.380 | 0.274 |

The topics *poruba, tram, line* and *snow, snowing, winter* have a similar mean sentiment value. However, the lower level of both positive and negative sentiment suggests that people agree more on the sentiment toward the former. Similarly, the topic *hydrogen, technology, h2* is perceived rather positively on average, but the high level of negative opinion can trigger further investigation. This information could not be extracted if only an arithmetic mean or other simple aggregation was used.

An interesting comparison can be made between topics *hydrogen, technology, h2* and *planetarium, transformation, night*. According to the mean sentiment, the former is perceived more positively than the latter. However, the higher level of positive sentiment and the lower level of negative sentiment for the latter contradict this conclusion. This inconsistency might be caused by the asymmetry of the positive and negative opinion fuzzy sets (Fig. 4.5).

In summary, the partial results presented in Table 5.7 suggest that the authorities should prioritize and allocate more funds to parking, car traffic, and healthcare. At the same time, they should investigate topics such as the use of hydrogen more deeply.

### 5.2.4 Tweets from Brno

The framework detected 98 distinct topics in tweets published from Brno. This is much less than the number of topics identified in tweets from Ostrava. It is possible to speculate that the discussions in Brno are somewhat less diverse.

TFNs representing eight of the most discussed topics are shown in Fig. 5.5. The topic with the highest mean sentiment is *fun, sport, plastic, time*. At the same time, this topic seems to be highly controversial. The topic perceived most negatively is *russian, russia, ukraine, ukrainian*. The narrow support for TFN representing this topic also signals

a relative unity of sentiment among different people. Considering the ongoing Russian invasion of Ukraine, this finding is not surprising.

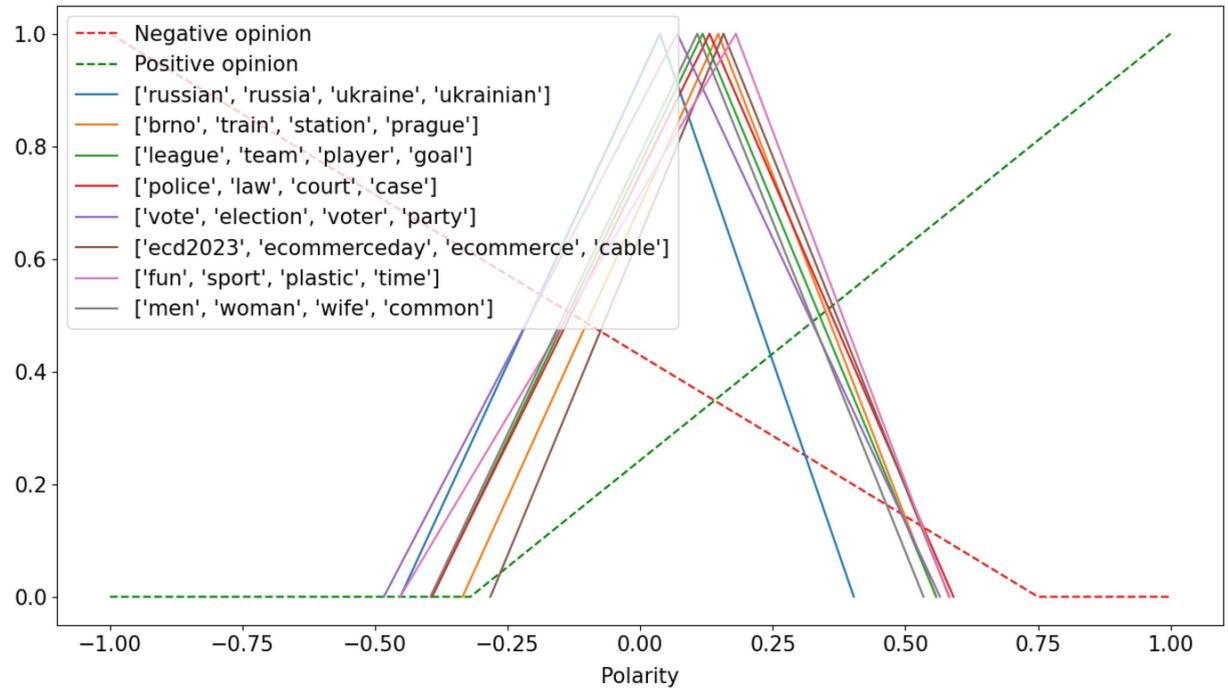


***Fig. 5.5*** TFNs representing aggregated sentiment toward eight most discussed topics detected in tweets published from Brno ($s = 1.5$)

When comparing the most discussed topics in the two cities, it is possible to see that people in both cities are interested in the Russian invasion of Ukraine, sports, as well as issues related to police and law. However, interest in other topics relevant to municipal decision-making such as education, healthcare, or construction seems to be much higher in Ostrava. One could hypothesize that citizens of Ostrava are more engaged or that these services work better in Brno, and there is less need to discuss them on social media.

The second hypothesis can be examined by comparing the topics that occur in both Tables 5.7 and 5.8. Healthcare is perceived much better in Brno than in Ostrava, but the sentiment toward construction and education is significantly more negative. These topics are the two most

negatively perceived in Table 5.8. The negative perception of education is rather surprising, as it goes against the popular idea that the universities in Brno are of high quality, at least compared to other universities in Czechia. A more positive outlook on parking compared to Ostrava is also rather unexpected. Given the larger population, it would be reasonable to expect that parking in Brno would be more problematic.

***Table 5.8*** Mean sentiment and scaled level of positive and negative opinion for several topics relevant to municipal decision-making in Brno

| Topic | Mean sentiment | Level positive | Level negative |
|---|---|---|---|
| **Brno, train, station** | 0.147 | 0.555 | 0.406 |
| **Car, parking, parked** | 0.119 | 0.439 | 0.430 |
| **Doctor, medicine, hospital** | 0.146 | 0.721 | 0.534 |
| **School, university, education** | 0.079 | 0.275 | 0.720 |
| **Inflation, rent, price** | 0.104 | 0.366 | 0.510 |
| **Tax, property, municipality** | 0.145 | 0.648 | 0.440 |
| **Fuel, scarcity, energy** | 0.127 | 0.511 | 0.489 |
| **Salary, employer, pay** | 0.123 | 0.503 | 0.438 |
| **Read, text, study** | 0.120 | 0.520 | 0.561 |
| **Building, construction, build** | 0.065 | 0.229 | 0.657 |

In general, the topics in Table 5.8 seem to be more controversial than the topics in Table 5.7, as there are many topics with high levels of positive and negative sentiment. Even topics such as *doctor, medicine, hospital* that have a relatively high mean sentiment still have a high level of negative opinion. However, it is important to note that the positive and negative opinion fuzzy sets used to calculate the levels of positive and negative opinions were estimated only from sentiment evaluations of tweets

published from Ostrava by people living in Ostrava. It remains unclear whether these results can be generalized and used when evaluating sentiment in other cities.

### 5.2.5 Comparison with a Naive Method

The Ostrava dataset was also processed using a naive method similar to that used by commercially available tools. Instead of determining the sentiment polarity on a continuous scale, each tweet was labeled positive (1) or negative (−1). These labels were then used to calculate the weighted mean sentiment, as well as two measures of opinion diversity: weighted standard deviation and entropy. The entropy was used with the intuition that a balanced number of positive and negative tweets can mean high opinion diversity.

The results of the naive approach were then compared with the results provided by the framework, and the correlation between mean sentiment and sentiment diversity was examined. For the proposed framework, the sentiment diversity was calculated as the length of the support interval. Table 5.9 shows the Pearson correlation values.

***Table 5.9*** Pearson correlation between mean sentiment and diversity measures for the proposed framework and the naive approach to sentiment analysis

| | **Mean (framework)** | **Mean (naive)** |
|---|---|---|
| **Mean (framework)** | – | 0.841 |
| **Diversity (framework)** | −0.252 | – |
| **StdDev (naive)** | – | −0.974 |
| **Entropy (naive)** | – | −0.969 |

Two conclusions can be drawn from Table 5.9. First, the correlation between mean sentiments calculated by the two methods is high, so both approaches are comparable when

determining mean sentiment. Second, the correlation between mean sentiment and diversity measures is low for the framework and very high for the naive approach. This suggests that the framework treats mean sentiment and sentiment diversity separately. In contrast, the naive approach essentially produces a single variable that can be interpreted in different ways. Therefore, it can be argued that the proposed framework extracts more information from the same data and reduces the ambiguity of interpretation.

### 5.2.6 Likes and Shares

In addition to posting new content, many social media platforms allow their users to *like* and *share* (*retweet* in case of Twitter) content posted by someone else. The total number of likes or shares a given post gets could be used as a weight during the fuzzy aggregation process.

There is a large body of research on how people use likes on social networks. The work [7] studied how people think about posting and liking on image-heavy social networks and confirmed the intuitive idea that liking an image means that a person likes the image. The research presented in [5] indicates that many people perceive liking as a gift to the author. The study showed that liking was positively correlated with interpersonal generosity, public self-confidence, and age and negatively correlated with education. The gift-giving view is consistent with those who found that likes are reciprocally used to maintain relationships [1]. People may care more about who likes their content than about how many likes they receive in total [9]. Finally, the study [10] found that giving a like is associated with an activity in parts of the brain associated with reward.

Research on sharing and retweeting is sparser. The study [8] suggests that the act of retweeting is similar to

direct quotation. It is an ostensive act whose purpose is to bring the audience's attention toward something. An older article [6] suggests several triggers for retweeting: source trustworthiness, source expertise, source attractiveness, and also presence of multimedia content in the post. According to [2], users retweet content that they not only like but also agree with. At the same time, some users may retweet for more self-centered reasons, such as establishing their social network profile as an influential information hub.

One could form two intuitive hypotheses about the relationship between the number of likes, the number of tweets, and sentiment: (H1) Tweets that are liked are also more frequently retweeted, and (H2) tweets that are highly positive or highly negative are more likely to catch users' attention and hence have more likes and retweets. The validity of these hypotheses was examined by calculating the correlation between these variables in the tweet dataset published from Ostrava.

The results are shown in Table 5.10. Both the original sentiment polarity and its absolute value were considered. The absolute value disregards the direction of the sentiment and focuses only on its strength. In both cases, the correlation between likes and sentiment, and between retweets and sentiment, is very low and not statistically significant, suggesting that hypothesis (H2) is rejected. In contrast, the correlation between the number of likes and the number of retweets is relatively high. This supports the validity of hypothesis (H1).

***Table 5.10*** Pearson correlation between the number of likes, the number of retweets, the sentiment polarity, and the absolute sentiment polarity for the dataset of tweets published from Ostrava

| | Likes | Retweets | Sentiment | Absolute sentiment |
|---|---|---|---|---|
| **Likes** | | ***0.709 | 0.009 | 0.011 |

| | Likes | Retweets | Sentiment | Absolute sentiment |
|---|---|---|---|---|
| **Retweets** | | | −0.012 | 0.017 |
| **Sentiment** | | | | 0.512 |

Both the number of likes and the number of retweets could be used as a weight when aggregating sentiment toward a topic into a TFN. Due to their high correlation, only one of these measures should be entered in the weight calculation (unless their interaction is properly handled). Based on the previously mentioned idea that retweets express both a positive view and an agreement, retweets were further examined as a possible weight coefficient for sentiment aggregation. Two variants were considered: the original number of retweets and its square root. The intuition behind the square root transformation is the potential existence of a positive feedback loop: The more retweets a Twitter post has, the easier it is for it to get even more retweets. The square root mitigates the effect of this feedback loop by creating a sublinear relationship between the number of retweets and the weight of the tweet. Since the probability of a tweet belonging to a given topic was already used as weight, the new overall weight is calculated by multiplying these two factors:

$$w_{post,topic} = p(post \in topic)(n_{retweets} + 1) \tag{5.1}$$

or

$$w_{post,topic} = p(post \in topic)(\sqrt{n_{retweets} + 1}). \tag{5.2}$$

Adding 1 to the number of retweets prevents posts with no retweets from having a weight of 0 and hence from being ignored during aggregation. The post weights for each topic are normalized to sum up to 1.

The effects of using the number of retweets as a weight for six selected topics are shown in Fig. 5.6. In some cases, such as the topic *czech, republic, czech republic, prague*, there is little to no difference. In the case of the topic *ukraine, russian, putin, russia*, it is possible to observe a shift of the TFN core. Finally, in the case of the topic

*doctor, hospital, patient, nurse* or *sport, home, carousel, year,* a wider and more asymmetric support can be observed.

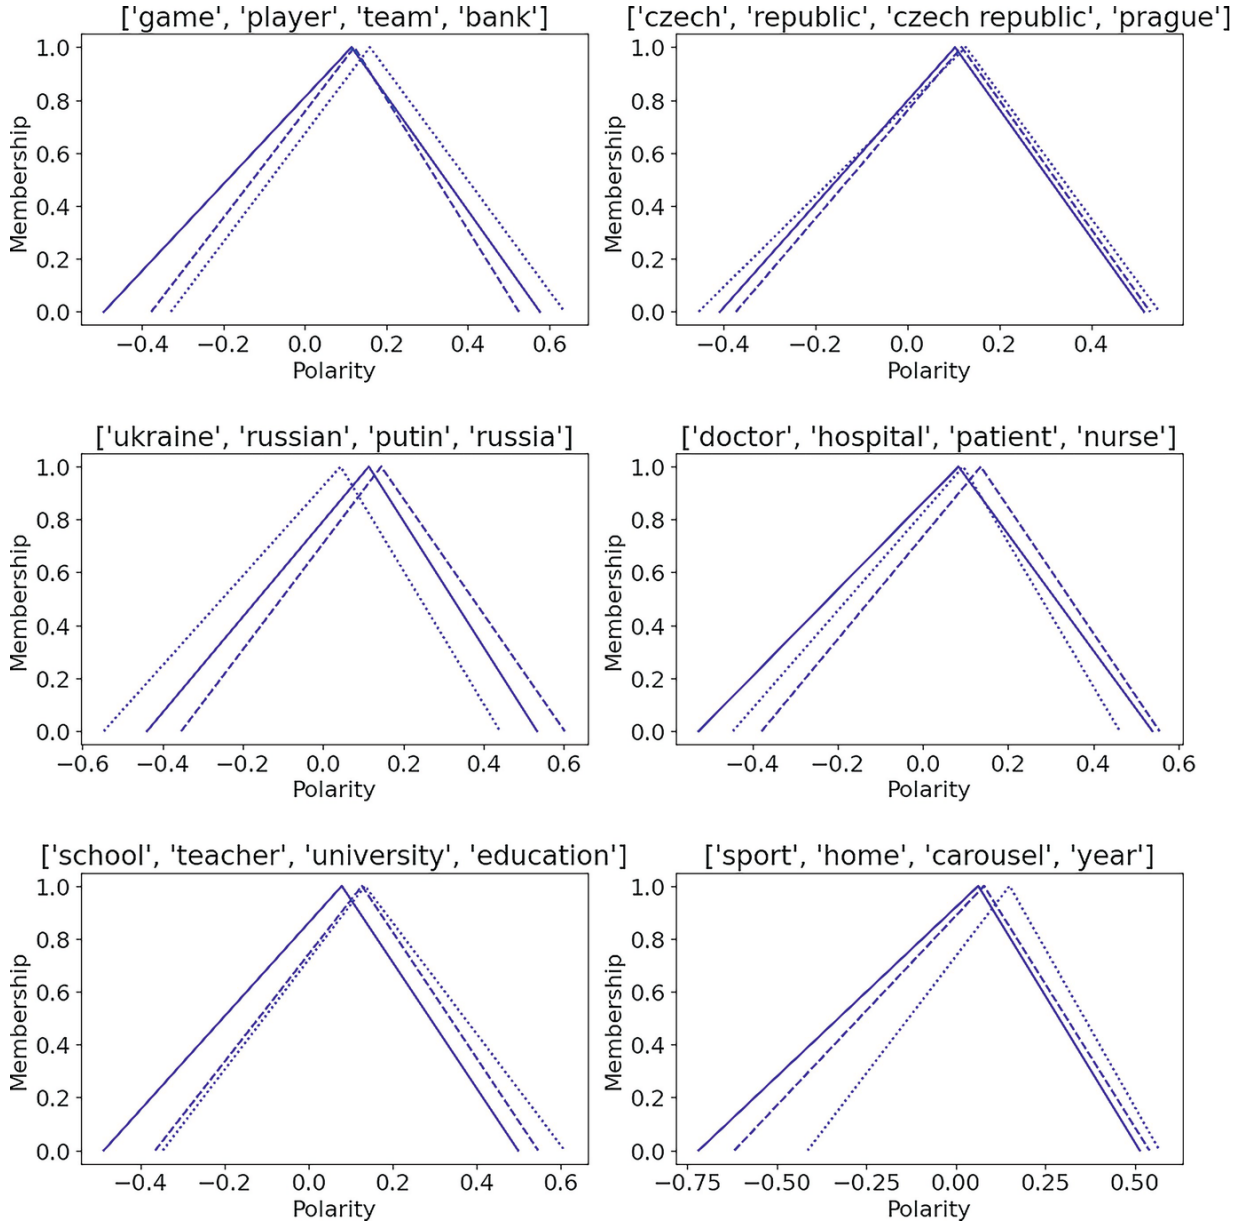


***Fig. 5.6*** Comparison of TNFs with a number of retweets as weights. Full line: number of retweets as weight; dashed line: square root of the number of retweets as weight; and dotted line number of retweets ignored

It is not possible to say that any of the three TFN variants is better than the others. One could argue that a broader and more asymmetric support helps to distinguish

between topics. This property might be beneficial in some use cases, such as when a multiple criteria decision-making method is applied to the data. In general, the use of retweets as weights should be user decision.

### 5.2.7 Web Application

For the framework to be practically useful, it has to be presented in a user-friendly form. Therefore, a prototype of a simple web application was developed to present the results of the framework. The application is named "OpinioNet." It allows the user to switch between multiple cities (currently only Brno and Ostrava) and provides both a comparison of different topics and a detailed analysis of each topic in isolation. At the time of writing, the application can be accessed at opinionet.swarm.svana.name.

The main dashboard consists of two parts: a sentiment overview chart shown in Fig. 5.7 and word clouds for positive and negative tweets shown in Fig. 5.8. The sentiment overview chart visualizes several variables related to the topics most discussed in the selected city. The green and red bars represent the overall level of positive and negative opinion for a given topic. The yellow square carries information about the mean sentiment conveyed by its placement and the discussion volume conveyed by its size.

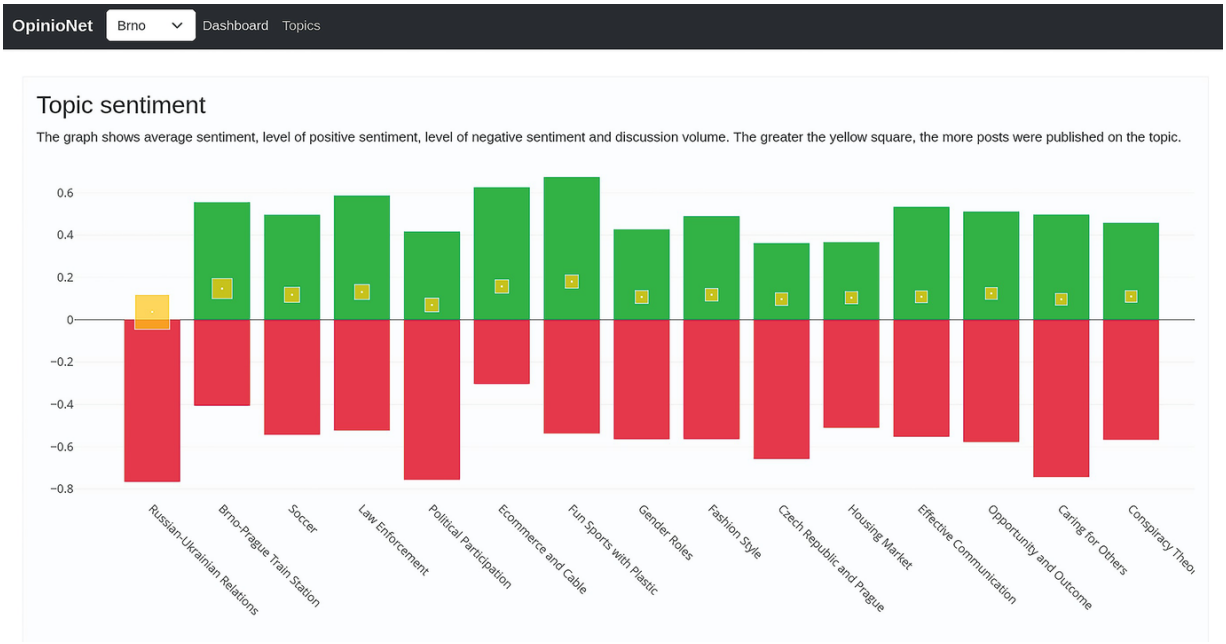


***Fig. 5.7*** OpinioNet: sentiment overview

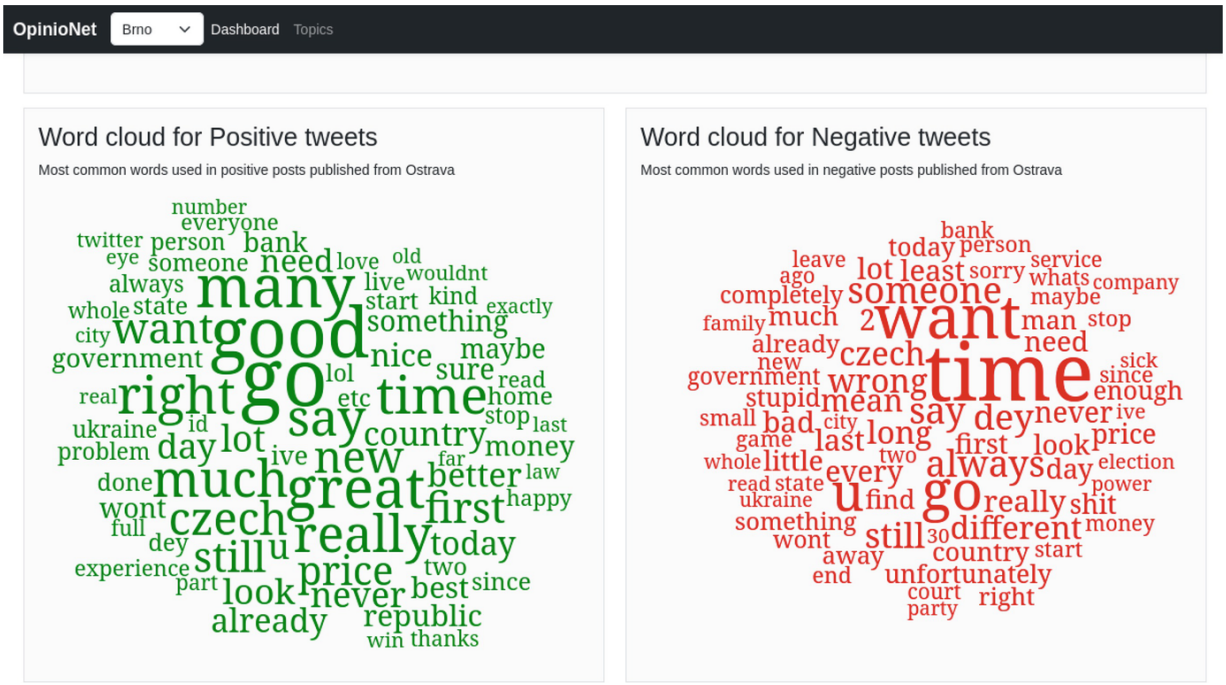


***Fig. 5.8*** OpinioNet: word clouds of most common words in positive and negative tweets

The word clouds display words that occur most frequently in positive and negative tweets. The size of each word represents its relative frequency. Words that occur equally frequently in both positive and negative tweets are

omitted under the assumption that these words do not carry useful information.

The user can also explore a full list of topics by clicking on the *Topics* link in the top panel. A screenshot of this list is shown in Fig. 5.9. The figure shows the mean sentiment, the levels of positive and negative sentiments, the level of controversy, and the volume of discussion for each topic. In addition to showing a number, the relative values of these measures are visualized with simple bars in the background. The top five values in each column are colored green, the bottom five values are colored red, and all other values are colored blue. By default, the topics in the table are sorted by discussion volume, but the user can sort the table by any other measure by clicking on a corresponding column name.

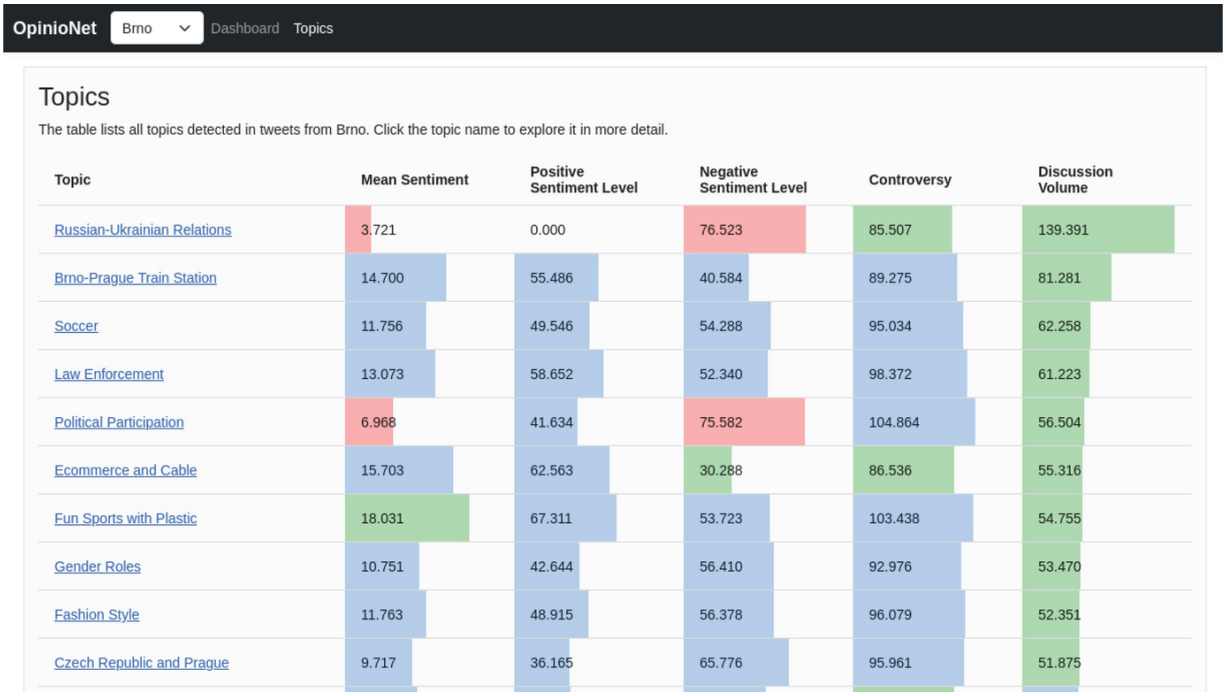
OpinioNet
Brno
Dashboard
Topics
Topics
The table lists all topics detected in tweets from Brno. Click the topic name to explore it in more detail.
Topic
Mean Sentiment
Positive Sentiment Level
Negative Sentiment Level
Controversy
Discussion Volume
Russian-Ukrainian Relations 3.721 0.000 76.523 85.507 139.391
Brno-Prague Train Station 14.700 55.486 40.584 89.275 81.281
Soccer 11.756 49.546 54.288 95.034 62.258
Law Enforcement 13.073 58.652 52.340 98.372 61.223
Political Participation 6.968 41.634 75.582 104.864 56.504
Ecommerce and Cable 15.703 62.563 30.288 86.536 55.316
Fun Sports with Plastic 18.031 67.311 53.723 103.438 54.755
Gender Roles 10.751 42.644 56.410 92.976 53.470
Fashion Style 11.763 48.915 56.378 96.079 52.351
Czech Republic and Prague 9.717 36.165 65.776 95.961 51.875

***Fig. 5.9*** OpinioNet: list of all identified topics

Finally, the user can see a detailed analysis of each topic by clicking on its name in the table. The topic details screen shows the values of all basic variables for the topic

in the form of a table and in the form of a radar chart. Moreover, the user can select a different topic and add it to the radar chart for comparison, as illustrated in Fig. 5.10.

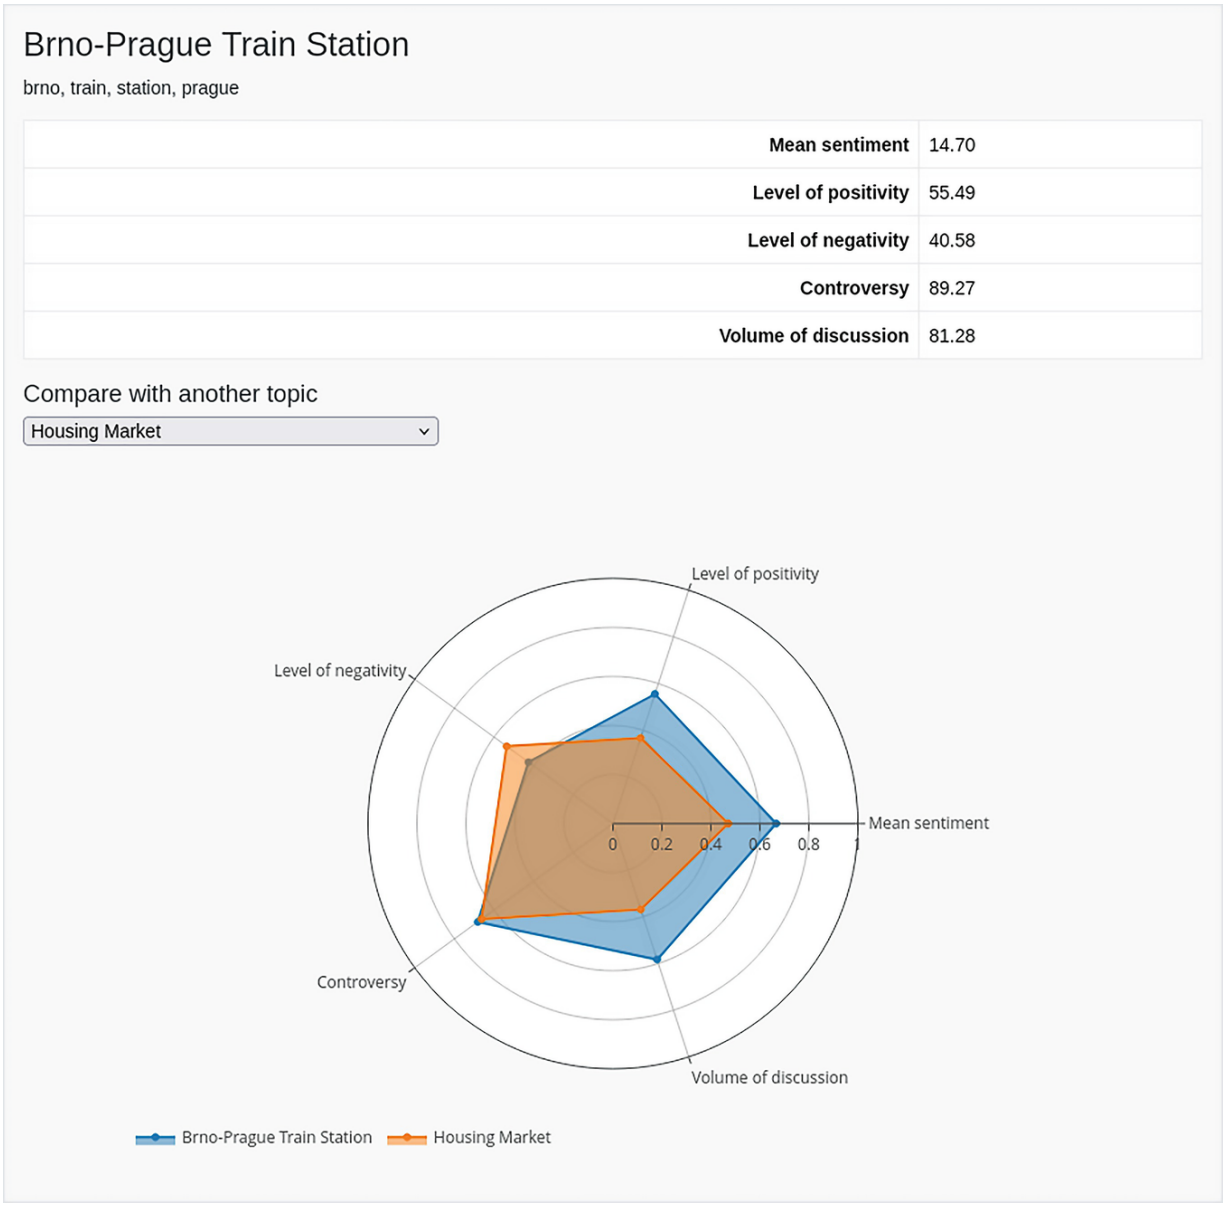


***Fig. 5.10*** OpinioNet: comparison of two topics

Instead of simply showing the most descriptive words, the application gives each topic an easily understandable title. This feature utilizes the capabilities of the GPT large language model, which is asked to generate a topic title with the following prompt:

> Suggest a short, human-readable name for a topic characterized by the following words and store it in the database. The topic was extracted using a topic modeling method:
>
> {topic_words}

The application also uses GPT to generate a summary of positive and negative tweets for each topic. Given that most of the tweets in the dataset are in Czech, the summaries also use this language. It was assumed that this choice minimized the loss of information due to translation. For example, the following summary of positive tweets was generated for the topic “Brno-Prague Train Station” shown in Fig. 5.10:

> Lidé z Brna vnímají na téma “Brno-Prague Train Station” pozitivně především praktičnost a pohodlí spojení mezi Brnem a Prahou. Ocenění si získává možnost cestování vlakem, která je považována za rychlou a efektivní, a to i v porovnání s jinými dopravními prostředky. Vlakové spojení je využíváno jak pro pracovní tak pro rekreační účely, a to i mezinárodně. Cestující si také cení možnosti pohodlného přestupu na regionální vlaky a dobré dostupnosti veřejné dopravy v oblasti hlavního nádraží. Pozitivně je vnímána i architektura a historický význam nádražní budovy, která je součástí městského prostředí a kultury.

The summary roughly translates to:

> People from Brno perceive the topic “Brno-Prague Train Station” positively, especially the practicality and convenience of the connection between Brno and Prague. The option of traveling by train, which is considered fast and efficient, is also appreciated in

comparison to other means of transport. The train connection is used for both work and leisure purposes, including internationally. Passengers also value the possibility of convenient transfers to regional trains and the good availability of public transport in the area of the main station. The architecture and historical significance of the station building, which is part of the urban environment and culture, is also viewed positively.

The following prompt is used to generate such summaries with GPT:

> Podívej se na následující tweety a napiš krátký odstavec v češtině, který sumarizuje, co přesně lidé z města {city} vnímají na tématu “{topic_name}” sentiment. Nezmiňuj ale aspekty, na které lidé nahlíží {opposite_sentiment}:
>
> TWEETY:
> {list_of_tweets}

which again roughly translates to:

> Have a look at the following tweets and write a short paragraph in Czech summarizing what exactly people from the city of {city} perceive {sentiment} on the topic “{topic_name}”. Do not mention aspects that people perceive {opposite_sentiment}:
>
> TWEETS:
> {list_of_tweets}

The prompt should prevent GPT from including negative aspects in the positive summary and vice versa to clearly separate the two.

Automated summaries can be very useful for getting detailed information on a given topic and deriving steps to improve the situation. As already discussed, such an

understanding is especially important in the case of highly controversial topics.

On the technical level, the application consists of two parts: a back end and a front end. The back end is responsible for providing all data for visualization. It is implemented in the Python programming language and utilizes the FastAPI[2] framework to provide an API that the front end can use to access the data. The back-end part also includes a series of preprocessing scripts responsible for transforming the data into a form suitable for visualization or generating topic titles with GPT. It is also worth mentioning that the back end performs an on-demand summary generation. Since GPT is a paid service when integrated into another application, in order to lower the costs, a summary is generated when a user visits the detail page of a given topic for the first time. This usually takes approximately 20–40 seconds. After a summary is generated, it is stored in a database. When a user visits a topic detail page for the second time, this stored summary is used and shown almost immediately.

The front end is written in TypeScript and uses the React[3] and Bootstrap[4] libraries as its backbone. It implements the actual user interface. The fundamental idea behind how the front end works is rather simple. When visiting a certain subsection, an HTTP request is performed to get the required data from the back end. Then it uses these data to render various visualizations.

The front end also utilizes libraries for various types of visualization. Charts are built on top of PlotlyJS,[5] word clouds on top of the react-d3-cloud library,[6] and the list of topics on top of the React Data React Data Table Component,[7] which provides functionality such as number formatting or sorting.

It should be noted that the application is only a prototype and that its functionality should be modified or

extended depending on the actual needs of municipal decision-makers or other types of users. Alternatively, it could allow for some degree of customization so that one single application can work well for different user groups.

---

## 5.3 Experiments with the Sentiment140 Dataset

The main intended use case of the proposed framework is the support for municipal decision-making. However, we also suggest that it can be generalized to other use cases. To demonstrate this, the framework is applied to a sample drawn from a well-known Sentiment140 dataset available in many online data repositories, such as Kaggle.[8] In this work, 50,000 randomly selected tweets from this dataset were processed in the same way as the set of tweets from Ostrava, the only exception being the omission of translation, since tweets in the Sentiment140 dataset are already in English.

Figure 5.11 shows the TFNs of eight topics most discussed in the Sentiment140 sample. The list contains topic that one would expect to see being discussed by an average social media user. Topics such as *school, exam, class, math* or *sleep, bed, tired, night* are perceived rather negatively, while the topic *songs, music, listening, album* is perceived relatively positively. Again, these results match intuitive expectations. It can also be observed that the topic *eating, eat, lunch, food* is perceived not only more positively than, for example, *school, exam, class, math* but also has a higher degree of diversity expressed by its wider support.

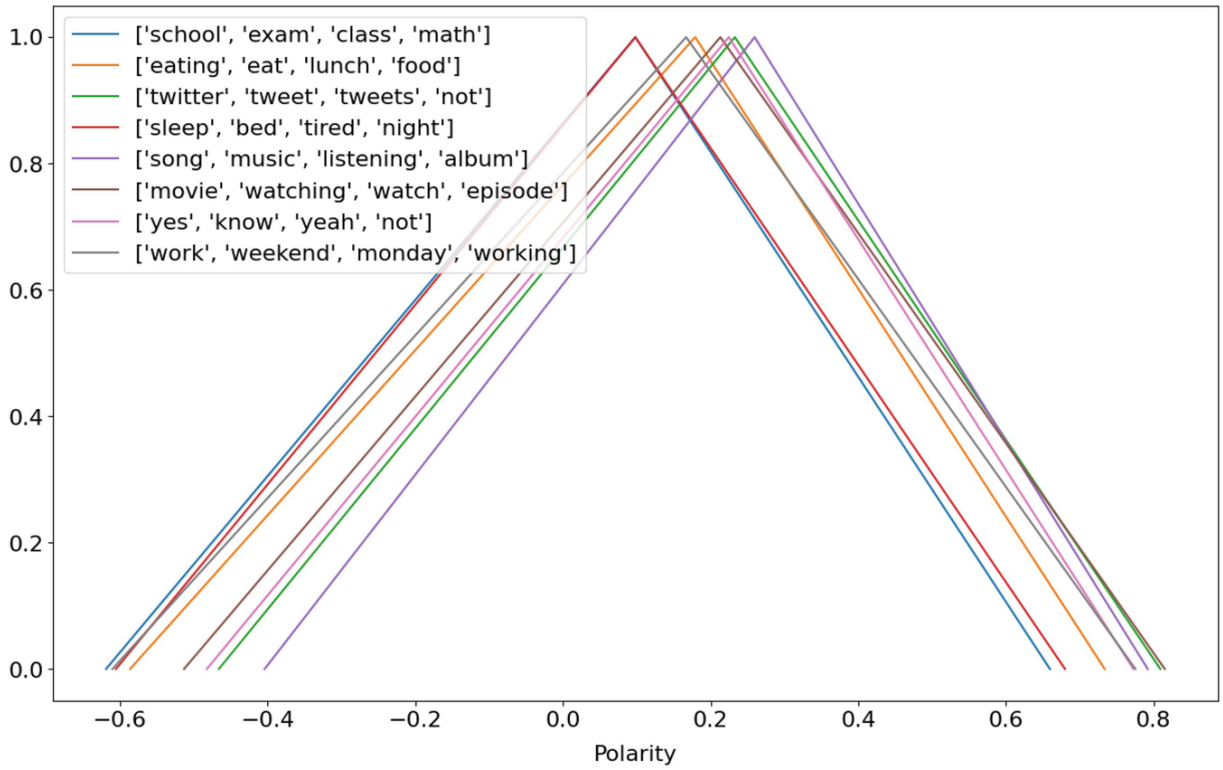


***Fig. 5.11*** TFNs representing aggregated sentiment toward selected topics detected in tweets from the Sentiment140 dataset ($s = 1.5$)

The same conclusions can be drawn from Table 5.11. For example, the topic *song, musing, listening* has not only the highest mean sentiment but also the highest level of positive opinion and the lowest level of negative opinion. In contrast, the topic *school, exam, class* has the lowest mean sentiment, the lowest level of positive opinion, and the highest level of negative opinion. Differences in the levels of positive and negative opinions between topics *school, exam, class* and *eating, eat, lunch* also signal more diversity for the latter.

***Table 5.11*** Mean sentiment and scaled level of positive and negative opinions for most discussed topics extracted from the Sentiment140 dataset

| Topic | Mean sent. | Level. pos. | Level. neg. |
|---|---|---|---|
| **School, exam, class** | 0.098 | 0.346 | 0.514 |
| **Eating, eat, lunch** | 0.179 | 0.550 | 0.398 |

| Topic | Mean sent. | Level. pos. | Level. neg. |
|---|---|---|---|
| **Twitter, tweet, tweets** | 0.232 | 0.713 | 0.228 |
| **Sleep, bed, tired** | 0.097 | 0.371 | 0.503 |
| **Song, music, listening** | 0.259 | 0.737 | 0.138 |
| **Movie, watching, watch** | 0.213 | 0.689 | 0.294 |
| **Yes, know, yeah** | 0.225 | 0.661 | 0.253 |
| **Work, weekend, Monday** | 0.167 | 0.578 | 0.433 |

The levels of positive and negative opinions in Table 5.11 were calculated using the fuzzy sets introduced in Sect. 4.5. However, for the reasons discussed in Sect. 6.2, this might not be the most appropriate approach. Therefore, the result presented in Table 5.11 should be treated with caution.

### 5.3.1 Comparison with a Naive Approach

To quantify the benefits of the proposed framework, both datasets presented in Sect. 5.2 were also processed by a naive method similar to those used by commercially available tools. Instead of determining the polarity on a continuous scale, each tweet was labeled as positive (1) or negative ($-1$). These labels were then used to calculate the weighted mean sentiment for each topic. Moreover, the diversity of opinion toward each topic was quantified using two indicators: weighted standard deviation and entropy. Entropy was used with the intuition that a balanced number of positive and negative tweets signals high opinion diversity and maximizes entropy at the same time.

The results of the naive approach were then compared with the results provided by the framework. More specifically, the correlation between mean sentiment and measures of sentiment diversity is analyzed. For the proposed framework, the sentiment diversity of each topic was calculated as the length of the support interval of its TFN. Tables 5.12 and 5.13 show the Pearson correlation

values for the Ostrava dataset and the Sentiment140 dataset, respectively.

***Table 5.12*** Pearson correlation between mean sentiment and diversity measures for the proposed framework and the naive approach to sentiment analysis (Ostrava dataset)

| | **Mean framework** | **Mean naive** |
|---|---|---|
| **Mean framework** | 1.0 | 0.841 |
| **Framework diversity** | −0.252 | — |
| **Naive StdDev** | — | −0.974 |
| **Naive entropy** | — | −0.969 |

***Table 5.13*** Pearson correlation between mean sentiment and diversity measures for the proposed framework and the naive approach to sentiment analysis (Sentiment140 dataset)

| | **Mean framework** | **Mean naive** |
|---|---|---|
| **Mean framework** | 1.0 | 0.968 |
| **Framework diversity** | −0.496 | — |
| **Naive StdDev** | — | −0.927 |
| **Naive entropy** | — | −0.921 |

According to both tables, the correlation between mean sentiments calculated by the two methods is high, so the two approaches are comparable when determining mean sentiment. However, the correlation between mean sentiment and diversity measures is low for the proposed framework and very high for the naive approach. This suggests that the framework treats mean sentiment and sentiment diversity separately, while the naive approach essentially produces a single variable that can be interpreted in different ways. It could be therefore argued that the proposed framework extracts more information from the same data and reduces the ambiguity of interpretation.

### 5.3.2 Processing Time

The time required to process the provided input was empirically evaluated for two parameters: the number of tweets and the minimum number of documents per topic. The second parameter has a direct impact on the number of detected topics. Performance was evaluated on a desktop PC with an i7-6700K CPU, NVIDIA GTX 980Ti GPU, and 32 GB of RAM. The system was running Ubuntu 22.04, and Python 3.11. BERTopic is able to utilize the GPU for a part of the topic detection process. All other parts of the framework were running in a single thread on a CPU.

Table 5.14 shows the running time for different numbers of tweets at the input for different parts of the computation while keeping the number of detected topics around 100. Most of the computation time is spent on topic modeling. The complexity of this component seems to be linear. This is reasonable, since the computational complexity is $O(n^2)$ or $O(n \log n)$, given $n$ data points and depending on the algorithm used (see [4]). The process of determining the levels of positive and negative sentiment runs in constant time, and all other parts of the framework scale linearly with the number of tweets.

***Table 5.14*** Running times (s) of different parts of the framework for different numbers of input tweets

| No. tweets | 25k | 50k | 100k | 200k | 400k |
|---|---|---|---|---|---|
| **Preprocessing** | 2.34 | 4.84 | 8.83 | 16.9 | 34.1 |
| **Topic modeling** | 43 | 93 | 172 | 440 | 1144 |
| **Sentiment analysis** | 4.37 | 8.7 | 17.4 | 35.6 | 69 |
| **Mean + semi-deviation** | 0.228 | 0.521 | 0.886 | 1.99 | 3.98 |
| **Pos./neg. levels** | 0.158 | 0.159 | 0.164 | 0.155 | 0.153 |

Table 5.15 shows the running time for different minimal numbers of tweets per topic with the number of input

tweets fixed at 50,000. As the minimal topic size increases and the number of detected topics decreases, the time required to perform topic modeling decreases slower than linearly. The time required to calculate the mean sentiment and its standard “semi-deviation” decreases approximately linearly. The time required for sentiment analysis and for determining the levels of positive and negative opinions appears constant.

***Table 5.15*** Running times (s) of different parts of the framework for different minimal topic sizes

| Min. tweets per topic | 16 | 32 | 64 | 128 |
|---|---|---|---|---|
| **No. topics** | 338 | 117 | 90 | 2 |
| **Topic modeling** | 321 | 128 | 86 | 68 |
| **Sentiment analysis** | 8.78 | 8.5 | 8.72 | 8.38 |
| **Mean + semi-deviation** | 1.28 | 0.682 | 0.393 | 0.024 |
| **Pos./neg. levels** | 0.152 | 0.156 | 0.153 | 0.136 |

In practice, it should be possible to reduce the time required to perform the topic modeling. BERTopic allows the user to train the topic model only on a subset of the whole dataset and then use this trained model to also assign topics to previously unseen tweets.

# References


1. Au A (2021) Guanxi 2.0: the exchange of likes in social networking sites. Inform Commun Soc 24(13):1891–1906. https://doi.org/1080/1369118X.2020.1748091 [Crossref]
2. Anger I, Kittl C (2011) Measuring influence on Twitter. In: Proceedings of the 11th international conference on knowledge management and knowledge technologies, pp 1–4. https://doi.org/1145/2024288.2024326
3. Bramer M (2020) Principles of data mining. Springer, London [Crossref]
4. Campello RJGB, Moulavi D, Sander J (2013) Density-based clustering based on hierarchical density estimates. In: Pei J, Tseng V S, Cao L, Motoda H, Xu

G (eds) Advances in knowledge discovery and data mining. Springer, Berlin, Heidelberg, pp 160–172
[Crossref]

5. Hong C, Chen Z, Li C (2017) “Liking” and being “liked”: how are personality traits and demographics associated with giving and receiving “likes” on Facebook? Comput Hum Behav 68:292–299. https://doi.org/1016/j.chb.2016.11.048
[Crossref]

6. Liu Z, Liu L, Li H (2012) Determinants of information retweeting in microblogging. Internet research: electronic networking applications and policy. https://doi.org/1108/10662241211250980

7. Lowe-Calverley E, Grieve R (2018) Thumbs up: a thematic analysis of image-based posting and liking behaviour on social media. Telematics Inform 35(7):1900–1913. https://doi.org/1016/j.tele.2018.06.003
[Crossref]

8. Marsili N (2021) Retweeting: its linguistic and epistemic value. Synthese. https://doi.org/1007/s11229-020-02731-y

9. Scissors L, Burke M, Wengrovitz S (2016) What’s in a Like? Attitudes and behaviors around receiving likes on Facebook. In: Proceedings of the 19th ACM conference on computer-supported cooperative work & social computing. Association for Computing Machinery, New York, pp 1501–1510. https://doi.org/1145/2818048.2820066

10. Sherman LE, Hernandez LM, Greenfield PM, Dapretto M (2018) What the brain “Likes”: neural correlates of providing feedback on social media. Soc Cogn Affect Neurosci 13(7):699–707. https://doi.org/1093/scan/nsy051
[Crossref]

---

# Footnotes

1 https://www.deepl.com/translator

2 https://fastapi.tiangolo.com/

3 https://react.dev/

4 https://getbootstrap.com/

5 https://plotly.com/

6 https://github.com/Yoctol/react-d3-cloud

7 https://react-data-table-component.netlify.app/

8 https://www.kaggle.com/datasets/kazanova/sentiment140

# Part III
# Conclusions and Recommendations

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2025
M. Švaňa et al., *From Text to Understanding*, Fuzzy Management Methods
https://doi.org/10.1007/978-3-032-00129-0_6

# 6. Discussion and Recommendations

Miloš Švaňa[1] ✉, František Zapletal[1] and Miroslav Hudec[1]
(1) VSB – Technical University of Ostrava, Ostrava-Poruba, Czech Republic

## 6.1 Biases

The framework is intended to be used primarily with data from social networks. However, it is unclear whether the users of social networks properly represent the whole population of a given city. For example, one could hypothesize that certain demographic groups such as elderly citizens are underrepresented. Several users dominating the discussion represent a second source of distortion.

The issue of bias that comes from underrepresentation of some group can be addressed in multiple ways, depending on its exact source. First, it can be recommended to combine the framework with at least one traditional opinion collection method, such as a municipal referendum or at least a guided live discussion between citizens and municipal representatives. Or, a municipality could use this framework to analyze free-form responses in surveys, which could be more representative of the overall population.

Some sources of bias can be mitigated by modifying the methodology if necessary. For example, tweets produced by very active users can be assigned lower weights during aggregation. If it is known that some group of citizens (social, gender, age) is underrepresented, then it can be at least partially compensated for this deficiency by giving higher weight to these tweets. On the other hand, this would not guarantee that the provided opinion is not biased (and just reinforced).

The weights would also be possible to use to reflect the relevance of the topic for a user. Currently, there are many people who always tempt to comment everything, regardless of their reasonable interest in the topic. For example, if the sentiment related to the building of planned train station is provided by someone who lives close to this location and who travels by train a lot, the opinion is intuitively more relevant than the one of someone who visits the location rarely and does not travel by train. On the other hand, such person can observe something which is overlooked by commuter. This idea has been used by [5] in the model of general municipal decision-making. However, when using the data from social networks, it is impossible to get the necessary input data (at least from the X platform).

Another possible problem can be caused by trolls and bots which can quite easily contaminate the overall opinion of real and interested users. If some bot just replicates a single tweet many times, it could be technically easy to remove the copies (or the whole set of posts) from the evaluation. However, AI tools can currently quite well rephrase the same idea (with a similar sentiment), and the identification would be much harder or even impossible. Especially if the collection of opinions from the social network was announced to the public in advance, there would be quite a high risk of trolling.

In terms of sentiment, one can expect a skewed distribution. Opinions with extreme sentiment will be discussed more often, reposted more often, and will inflame more emotions. The complainers and screamers will always be seen (and heard) more. That is why the results need to be treated with caution. Then, it also depends on what policy-makers expect from the model. The idea of an alternative to a traditional referendum would be somewhat naive. On the other hand, today social networks serve to some extent as a mirror that reflects the broader mood in society. Moreover, in addition to reflecting current sentiment on social networks, it is also often created and amplified there.

The last issue to be mentioned here is also related with the special (to some extent ephemeral) nature of the sentiment, and it is time dependency. The results of the analysis can be heavily impacted by the time period used for the input data. Here, the purpose of modeling is driving. It is possible to explore short-term sentiment caused by some particular shock or event (some kind of hysteria).

---

## 6.2 Determining the Shapes of Positive and Negative Opinions

The results of the framework depend on the calibration of the fuzzy sets “positive opinion” and “negative opinion” (Fig. 4.5). This calibration was performed using subjective ratings of polarity and sentiment strength extracted from a survey, see Sect. 4.5. This approach is flexible and straightforward, but it also has limitations in terms of reliability and generalization. The following issues should be taken into consideration.

- **Language dependency** [4]—different languages require different training data. Even if a translator of the highest quality is used, there is still no guarantee that the

translation will not change the strength or even the polarity of the sentiment of the text. Therefore, it can be clearly recommended to perform the calibration on the data in the same language as the corpus being processed.

- **Culture dependency** [1]—people of different nationalities may have different perceptions of sentiment levels. For example, the authors of [1] point to significant differences in the expression of emotion between Americans, Russians, and Japanese. The recommendation is therefore to calibrate the model based on the opinions of people of the same nationality as those who provided the corpus being processed. In the case of a multicultural corpus, either differentiation (calibration for each group separately) or at least a particularly thorough sensitivity analysis of the results obtained may be recommended.
- **Domain dependency** [4]—the degree of perceived sentiment is inextricably linked to emotions, which can vary considerably depending on the subject of the text. Let us have the following tweet: “According to statistics, dozens of frogs end up under the wheels of cars completely unnecessarily. It is so sad.” There is probably no dispute that the polarity of sentiment will be negative. However, its intensity will depend on the strong relationship the respondent has with nature and animals. If the same tweet were to be altered by replacing the word “frogs” with “children,” with all due respect to frogs, a much stronger emotion and a more strongly negative sentiment can be expected after this change. If the corpus dealt with high crime rates in the city and the calibration was done with topics such as planting flowers in public space, the results could be heavily biased. It is typical for urban decision-making that a wide range of topics of varying importance are addressed, and therefore it is virtually impossible to avoid this effect. A recommendation to at least partially mitigate the effect is

to calibrate using texts that at least approximate the distribution of topics in the target corpora.
- **Time dependency** [2]—different topics may elicit different strengths of reaction at different times. The calibration settings should therefore be updated before each investigation to minimize the time between calibration and source text generation.

---

## 6.3 Other Methodological Issues

The presented framework is only a prototype, and, as such, it suffers from several methodological issues that should be addressed in future research.

- Determining whether a given document is positive or negative is a rather simplistic form of sentiment analysis. It should be possible to distinguish between **different types of emotion**. For example, both anger and fear belong to the set of negative emotions. However, distinguishing the two can be beneficial in many situations, including municipal decision-making. Fear might imply that a certain project is not explained properly. On the other hand, anger might signal that a project is in conflict with the value system of people living in the city.
- The framework does not use information about the **amount of likes or shares**, which could be interpreted as an agreeable position, and hence influence the weight of a given post during aggregation.
- The framework can identify **topics irrelevant to a given use case**. Improving the topic modeling component to evaluate the usefulness of a topic for intended use might improve the overall user experience.
- The framework considers each post as an independent entity, although many social networks organize posts into conversational threads. Analyzing posts **in relation to**

**other posts** could again lead to better results and understanding. Moreover, many posts contain not only textual content but also **images or videos**, which can carry information about sentiment and the topic of the post.
- The framework cannot distinguish between multiple **locations and moments in time**. Both factors can be useful in practice. Distinguishing between locations can help individuals with their own decision-making, for example, when deciding which city or part of a city should move. Analyzing how sentiment changes over time can help municipalities evaluate the success of different policies or projects.

---

## 6.4 Applications

Finally, the framework can be applied to other types of data and use cases. An alternative source of data has been already mentioned: free-form responses in surveys. But returning to the realm of social networks, it is not difficult to imagine that the framework can be applied to data from other services such as Reddit, which already has discussion groups for many cities or countries.

The framework can also be used to analyze user reviews on platforms such as Google Maps, notes from town hall meetings, focus group sessions, or interviews. However, some methodological modifications will probably be required, since the documents in these scenarios can be much longer and cover multiple topics with different sentiment expressed in each of them.

For the framework to be useful in practice, it should also be transformed into a user-friendly application that provides intuitive visualizations. It should also allow the user to further analyze a specific topic, for example, by displaying a short summary of associated social network

posts. Such functionality could be implemented with tools such as generative language models.

The topic for future research and application is integrating the results from the proposed framework with concept linguistic summaries or short-quantified sentences of natural language and interpreting their validities on maps. For example, a sentence like *the majority of tweets related to pollution are rather negative* has validity in the unit interval which can be straightforwardly converted into color hues for regions on a map. The application for the medical domain is proposed in [3]. The proposed framework has adopted fuzzy sets and fuzzy logic, so further improvement might use this background.

# References


1. Engelmann JB, Pogosyan M (2013) Emotion perception across cultures: the role of cognitive mechanisms. Front Psychol 4:118. https://doi.org/10.3389/fpsyg.2013.00118 [Crossref]
2. Huang F, Yuan C, Bi Y Lu, J (2020) Exploiting long-term dependency for topic sentiment analysis. IEEE Access 8:221963–221974
3. Hudec M, Malovcová K, Trumic R, Rakovská E (2022) A new quality measure and visualization of the short-quantified sentences of natural language on maps—a case on COVID-19 data. Informatica 33(2):321–342 [Crossref]
4. Ligthart A, Catal C, Tekinerdogan B (2021) Systematic reviews in sentiment analysis: a tertiary study. Artif Intell Revi 54:4997–5053 [Crossref]
5. Zapletal F, Hudec M, Švaňa M, Němec R (2023) Three-level model for opinion aggregation under hesitance. Soft Comput. https://doi.org/1007/s00500-023-07853-2

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2025
M. Švaňa et al., *From Text to Understanding*, Fuzzy Management Methods
https://doi.org/10.1007/978-3-032-00129-0_7

# 7. Conclusions

Miloš Švaňa[1] ✉, František Zapletal[1] and Miroslav Hudec[1]
(1) VSB – Technical University of Ostrava, Ostrava-Poruba, Czech Republic

## 7.1 Ethics and Sustainability

The use of the proposed framework as a decision support tool in the real world poses several ethical and sustainability challenges that should be taken into account. Three of the most important challenges are privacy concerns and consent, bias, and access to information.

**Privacy Concerns and Consent** The handling of private data by social media companies is a controversial topic. They are often used for morally questionable purposes, such as targeted advertising. Moreover, users have little control over how their data are processed and stored when they publish them on social networks. The framework can extract a lot of information that can be considered private by many social networks users. For instance, the Twitter API provides information about users' names or location. Additional information about users can be extracted from their posts, as well as from their likes or retweets. Although users share this information voluntarily, they may not have explicitly consented to their data being used by municipalities. The issue of privacy and context can be addressed in multiple ways. First, collected information can

be reduced to the bare minimum required for the framework to function. Second, users of the framework can have limited access to fine-grained information. For example, they could be prevented from accessing individual tweets. Third, users can explicitly opt out of data collection.

**Bias** As already mentioned, social media users may not accurately represent the entire population. Certain groups, such as elderly citizens, people with limited access to the Internet, or people not interested in sharing their private life details, can be underrepresented. At the same time, some users might dominate the discussion by producing a large number of posts. If the usage of the framework by a municipality becomes public knowledge, some users might try to intentionally shift the polarity of certain issues. The impacts of such biases could be mitigated in at least two ways: by combining the framework with additional sources of data and by implementing methods to handle users that dominate the conversation.

**Access to Information** It is unclear who should exactly have access to the results of the framework. Should it be restricted only to municipal workers and related personnel, or should it be granted to researches or even to the general public? Availability to the general public could inform each person's personal decisions, motivate them to participate by seeing that their opinions are being used, or introduce a new level of participation by discussing the results. At the same time, some people might object to a lower level of privacy or to the potential misuse of available information by malicious actors.

**Sustainability** The use of the framework itself should have a low impact on environment and climate, at least in terms of energy usage. Most components of the framework can run on relatively low-power devices, such as consumer-

level laptops. The only weakness is the use of large language models, for example, for generating tweet summaries. Compared to other types of software, large language models require significantly more computational resources to operate and therefore significantly more energy [1]. However, they are not a key component of the framework, and relatively simple techniques such as caching can be applied to minimize their use.

---

## 7.2 Future Research Opportunities

The proposed framework has several limitations. Part of them was introduced purposefully to limit the scope of the work, while the other were discovered during research and the framework design. These limitations are inspirations for future research opportunities.

First, the framework could be used to process other types of data. Instead of relying primarily on text, many, if not most, social network users prefer to post pictures and videos. Optical character recognition (OCR) and speech-to-text technologies could be used to extract linguistic data from these types of content. Furthermore, recently introduced multimodal language models, such as GPT-4,[1] could be used to generate textual descriptions of what is actually depicted in images and videos.

Second, it might be beneficial to consider the broader context of each post instead of treating it in isolation. Social media posts are often part of larger discussion threads as they respond to previous posts created by other users. Other posts in the discussion thread can provide information that improves topic detection and assignment. Information about the author of the post could also be considered. For example, it could be assumed that persons who publish. Because later in sentence is their large amounts of content on a certain topic understands it well

and that their opinion could have greater weight. On the other hand, such a person might tend to influence opinions in the opposite direction.

Third, many topics detected by the framework are not really relevant to municipal decision-making. Their presence can make the results harder to use and analyze by municipal decision-makers. Modern large language models or traditional linguistic summaries could be used to determine whether a given topic should be included in the final report. Moreover, the same approach could be applied even sooner in the process to filter out irrelevant social media posts. A potential implementation could involve comparing embeddings of expressions such as *city*, *municipality*, or *municipal decision making* with embeddings created from the topic description or the social network posts themselves.

Fourth, the results pointed out an inconsistency when ranking topics according to mean sentiment and according to the level of positive and negative opinion. At the surface level, this issue is caused by the asymmetry of the fuzzy sets, formalizing positive and negative opinions. However, it can be speculated that the underlying cause lies in the limitations of the sentiment analysis model. Therefore, alternative sentiment analysis methods should be explored. Creating a specialized sentiment analysis model designed specifically for municipal decision making could also be beneficial.

Fifth, there is a potential issue of social influence in relation to the concept of “wisdom of crowds.” The wisdom of crowds refers to the phenomenon that crowds can accurately estimate various quantities by averaging individual estimates. The independence of individual estimates is traditionally considered as one of the key requirements for the wisdom of crowd effect to take place. If this requirement is true, the framework could provide

biased information, as social media users are not prevented from seeing each other's content.

However, the literature on the topic presents mixed results. The work [6] used an analytic approach to show that social influence indeed has a negative effect on the outcome. But the authors also identified an exception: Social influence can help if the initial error is large. The study [5] similarly says that social influence is helpful if individuals lack prior knowledge. An older paper [7] describes a simulation in which individuals are represented as Brownian particles and concludes that the initial configuration is more important than the social influence. The study [2] reveals that the wisdom of crowds is robust and emphasizes the importance of differences in individual judgment. The impact of incorrect or misleading information is especially relevant in the context of social media. However, [4] suggests that crowd judgment is quite resistant to such issues.

Assuming that the problem of social influence actually exists, some researchers proposed solutions to mitigate it. For example, [3] proposes using information about sensitivity to social influence to detect individuals who are more resistant and focus on their estimates, and [8] illustrates how discussion in small groups and a consequent reevaluation of one's estimate can mitigate the problem.

Future research is required to find evidence for or against the presence of issues related to social influence in the framework. If such issues are present with a reasonable degree of certainty, methods for mitigating them should be developed.

Sixth, it was discovered that a single user can publish a significant portion of all posts in a given city. This person can have a disproportionate impact on the results. This problem should be addressed either directly by considering the number of posts per user during sentiment aggregation

or indirectly by considering other metrics, such as likes or retweets.

In addition to methodological improvements, the research presented in this work could also be extended by considering other applications of the proposed framework. The application space can be broadened in two ways: (1) data sources and (2) target user groups and domains.

In terms of data sources, the framework could be applied to data from other public social networks, such as Reddit. There are many Reddit communities (*subreddits*) associated with different cities and countries around the world. Compared to X (Twitter), Reddit places more emphasis on discussion between users. A typical interaction begins with a user asking a question, writing a longer post expressing their opinion, or sharing a news article. Other users can then respond to the original content or other users' responses, creating a tree of posts. It could be assumed that this mode of interaction leads to richer, more nuanced discussions.

The framework itself is rather universal and can be applied to any dataset that contains sentiment-expressing text documents discussing different topics. Examples include user reviews, comments on a news article, or free-form responses from surveys. This universality opens the way for the application of the framework in other domains. Data related to cities could be used not only for municipal but also for private decision-making. Businesses could use the data to help them decide where to open a new branch or factory. Individuals could make comparisons between cities when deciding where to move. Data could also help researchers studying fields such as urban planning, economics, sociology, politics, or international relationships.

If other data sources are considered, the framework can be a useful tool for businesses in general. It can help them

analyze the relationships their customers have with their products and services, analyze their weaknesses, and provide more benefits to the consumer.

---

## 7.3 Concluding Remarks

This book contributes a robust framework for sentiment analysis tailored to social media content, specifically aimed at supporting municipal decision-making in smart cities. By integrating various computational tools and sentiment models, the framework not only discerns the polarity of public opinion but also captures the diversity and complexity of sentiments expressed on digital platforms. All design choices and parameters' setting have been explored carefully by presenting a set of partial experiments to evaluate tools and libraries that perform different roles in the framework. Demonstrated through a real-world case study using tweets from the Czech Republic, the framework identifies pressing urban topics and evaluates the community's emotional responses. These insights equip decision- and policy-makers with actionable intelligence, enabling more responsive and inclusive urban governance. The framework is able to discover topics that were interesting to municipal decision-makers and assess the sentiment expressed about these topics. By aggregating the sentiment into a TFN and calculating the level of positive and negative opinions, it was able to recommend topics that should be prioritized or further investigated.

For citizens, this methodology represents a significant step toward participatory governance. It facilitates a digital feedback loop where residents' voices, conveyed through their online expressions, are systematically analyzed and potentially translated into policy priorities. This empowers individuals, making them active contributors to urban planning rather than passive recipients of top-down decisions. For society in general, the possible implications

are broader. The framework illustrates how big data and artificial intelligence can democratize public administration, allowing diverse communities to articulate concerns and aspirations. From a scientific standpoint, this research extends the current state of knowledge by marrying topic modeling, fuzzy logic, and sentiment quantification into a coherent analytical tool.

In essence, the research presented in this book not only enhances theoretical understanding but also offers practical tools for building smarter, more empathetic cities, where governance is data, informed, inclusive, and responsive to the nuanced emotional landscape of its residents. On the other hand, limiting factors were also identified that limit the practical application of the proposed framework in municipal decision-making. These include in particular the effects of various biases that give rise to questions such as:

- How to guarantee sufficient representativeness of the sample for all cohorts concerned?
- Are all opinions equally important and relevant for the final decision?
- How to face trolls and bots which can heavily distort the opinion when gathering data published in public space?
- How to deal with interdependence of posts/opinions such as belonging to some threads, shares, likes, and retweets?
- How to guarantee the time relevance of the results taking into account the hard to estimate perception of sentiment over time?
- The black-and-white perception simplifying sentiment perception (positive vs. negative) has been replaced within the framework by a “monochromatic” more accurate fuzzy approach (to some extent positive/negative). However, there are a number of strong and important sentiments that polarity cannot

always capture accurately enough (fear, tension, nostalgia, etc.).

Some of the mentioned issues are related to the current data unavailability, and some of them present challenges which can be the subject of the future research. As the most of other general frameworks based on mathematical models, this one also provides partial modularity and versatility, allowing users to customize it for a specific problem.

# References


1. Brown T, Mann B, Ryder N, Subbiah M, Kaplan J D, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A, Agarwal S, Herbert-Voss A, Krueger G, Henighan T, Child R, Ramesh A, Ziegler D, Wu J, Winter C, Hesse C, Chen M, Sigler E, Litwin M, Gray S, Chess B, Clark J, Berner C, McCandlish S, Radford A, Sutskever I, Amodei D (2020) Language models are few-shot learners. In: Larochelle H, Ranzato M, Hadsell R, Balcan M F, Lin H (eds) Advances in neural information processing systems. Curran Associates, pp 1877–1901
2. Davis-Stober C, Budescu D, Dana J, Broomell S (2014) When is a crowd wise? arXiv preprint arXiv:1406.7563
3. De Polavieja G, Madirolas G (2014) Wisdom of the confident: using social interactions to eliminate the bias in wisdom of the crowds. arXiv preprint arXiv:1406.7578
4. Jayles B, Escobedo R, Cezera S, Blanchet A, Kameda T, Sire C, Theraulaz G (2020) The impact of incorrect social information on collective wisdom in human groups. arXiv preprint arXiv:2003.06160
5. Jayles B, Kim H, Escobedo R, Cezera S, Blanchet A, Kameda T, Sire C, Theraulaz G (2017) How social information can improve estimation accuracy in human groups. Proc Natl Acad Sci 114(47):12620–12625. https://doi.org/1073/pnas.1703695114 [Crossref]
6. Mavrodiev P, Schweitzer F (2021) The ambiguous role of social influence on the wisdom of crowds: an analytic approach. Phys A: Stat Mech Appl 567:125624. https://doi.org/1016/j.physa.2020.125624 [Crossref]
7. Mavrodiev P, Tessone C J, Schweitzer F (2012) Effects of social influence on the wisdom of crowds. arXiv preprint arXiv:1204.3463

8. Navajas J, Niella T, Garbulsky G, Bahrami B, Sigman M (2017) Aggregated knowledge from a small number of debates outperforms the wisdom of large crowds. arXiv preprint arXiv:1703.00045

---

# Footnotes

1 https://openai.com/gpt-4

# Index

## Symbols

$\alpha$-cut 23, 28, 34

## A

Attention mechanism 44

## B

Backpropagation 44
BERTopic 52, 66
Bias 99, 105

## C

Calibration 100
Citizen participation 4, 9
City-as-a-machine 8
Classification 38, 41
Classification metrics 38
Clustering 39, 42
Complement 27
Computing with words 46
Core 23
Correlation 76, 78, 83, 84, 93
Crowds wisdom 107

## D

Data-driven decisions 8
Decoder 44
Deep learning NLP 47
Defuzzification 25
Density-Based Spatial Clustering of Applications with Noise (DBSCAN), 42
Dimensional reduction 40
Distance-based similarity measures 30

## E

EEmpowerment 9
Emotion 101
Encoder 44
Enta’s similarity 32
Entropy 83
Euclidean distance 31

## F

F1 score 65
Fuzzy aggregation 63
Fuzzy extensions 27, 29
Fuzzy logic 9, 22, 46
Fuzzy set core 23
Fuzzy sets 22, 46
Fuzzy set support 23

## G

Gregson’s similarity 32

## H

Hamming distance 30
Hashtag 74
HDBSCAN 43
Hierarchical clustering 39
High-level NLP 46

## I

Internet of Things (IoT) 1

## K

Koczy similarity 31

## L

Laplace smoothing 42
Latent Dirichlet Allocation (LDA) 48, 50, 65
Latent Semantic Analysis (LSA) 48
Lexica 54
Linguistic scale 33

Linguistic variable 32
Low-level NLP 46

## M

Machine Learning 37
Mean squared error 69
Membership function 23
Minkowski distance 31

## N

Naive Bayes 41, 54, 65
Natural language processing 46
Necessity measure 29
Negative opinion 68, 80, 91, 107
Neural network 43

## P

Perceptron 43
Performance 94
Pervasive computing 13
Positive opinion 68, 79, 91, 107
Possibility measure 29
Privacy 105

## R

Referendum 10
Regression 38
Regression metric 38
Reinforcement learning 40
Representative clustering 39
Restle’s similarity 32
Retrofitted smart cities 4
Retweet 85

## S

Santini similarity 31
Self-supervised learning 41

Semi-supervised learning 41
Sentiment analysis 29, 42, 54, 55, 63, 72, 73, 75107
Set theory 22
Similarity measure 30
Smart citizens 6
Smart city 1, 3, 55
smart economy 2
Smart environment 2
Smart-from-the-start 3
Smart government 3
Smart living 3
Smart mobility 2
Smart people 2
Social networks 9
Supervised learning 37
Supremum distance 31
Sustainability 106

## T

*T*-conorms 25
TextBlob 65, 68
*T*-norms 25
Topic coherence 50
Topic modeling 47, 49, 50, 63, 73, 75, 94
Topic model quality 49
Transformers 44, 47
Translation 100
Trapezoidal fuzzy set 23
Triangular fuzzy number 24, 66, 73
Triangular fuzzy set 81, 85
Tweet 74, 78, 84
Tweet sentiment 76

## U

Unsupervised learning 39
Urban development 5

## V

Variance 67
Visualization 90

## W

Weighted semivariance 67
Williams and Steele similarity 31
Word clouds 87
Word embeddings 47

## Z

Zadeh’s extension principle 27